

Europäisches Patentamt

European Patent Office

Office européen des brevets



(11) EP 1 229 442 A2

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
07.08.2002 Bulletin 2002/32

(51) Int Cl.7: G06F 9/46, G06F 17/60,
H04L 12/56

(21) Application number: 02250431.0

(22) Date of filing: 22.01.2002

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: 22.01.2001 US 263573 P
14.02.2001 US 268893 P
24.04.2001 US 286225 P
31.07.2001 US 308932 P

(71) Applicant: SUN MICROSYSTEMS, INC.
Palo Alto, California 94303 (US)

(72) Inventors:
• Traversat, Bernard A.
San Francisco, California 94109 (US)
• Slaughter, Gregory L.
Palo Alto, California 94306 (US)
• Saulpaugh
San Jose, California 95120 (US)
• Abdelaziz, Mohamed M.
Santa Clara, California 95051 (US)

• Duigou, Michael J.
Fremont, California 94555 (US)
• Pouyoul, Eric
San Francisco, California 94131 (US)
• Hugly, Jean-Christophe
Palo Alto, California 94303 (US)
• Yeager, William J.
Menlo Park, California 94025 (US)
• Pabla, Kuldipsingh
Santa Clara, California 95051 (US)
• Gong, Li
Los Altos, California 94024 (US)
• Joy, William N.
Aspen, Colorado 81611 (US)
• Clary, Michael J.
Monte Sereno, California 95030 (US)

(74) Representative: Harris, Ian Richard
D. Young & Co.,
21 New Fetter Lane
London EC4A 1DA (GB)

(54) Peer-to-peer computing architecture

(57) A system and method for providing an open network computing platform designed for peer-to-peer computing. The peer-to-peer platform may provide protocols for peer-to-peer services and applications that allow peers to discover each other, communicate with each other, and cooperate with each other to form peer groups. The protocols may include a peer membership protocol, a peer discovery protocol, a peer resolver protocol, a peer information protocol, a pipe binding protocol, and a peer endpoint protocol. Services and applications that participate in the protocols may be provided to deal with higher-level concepts. Advertisements may be used to publish peer resources. The peer-to-peer platform provides the ability to replicate information toward end users and may enable peers to find content that is closest to them. The peer-to-peer protocols and unique peer identifiers may allow peer nodes to move to different locations and access services and other content independent of network physical addresses.




FIG. 2

Description

BACKGROUND OF THE INVENTION5 1. Field of the Invention

[0001] This invention relates to peer-to-peer networking, and more particularly to a peer-to-peer network computing platform.

10 2. Description of the Related Art

[0002] The Internet has three valuable fundamental assets - information, bandwidth, and computing resources - all of which are vastly underutilized, partly due to the traditional client-server computing model. No single search engine or portal can locate and catalog the ever-increasing amount of information on the Web in a timely way. Moreover, a huge amount of information is transient and not subject to capture by techniques such as Web crawling. For example, research has estimated that the world produces two exabytes or about 2×10^{18} bytes of information every year, but only publishes about 300 terabytes or about 3×10^{12} bytes. In other words, for every megabyte of information produced, only one byte gets published. Moreover, Google claims that it searches about only 1.3×10^8 web pages. Thus, finding useful information in real time is increasingly difficult.

[0003] Although miles of new fiber have been installed, the new bandwidth gets little use if everyone goes to one site for content and to another site for auctions. Instead, hot spots just get hotter while cold pipes remain cold. This is partly why most people still feel the congestion over the Internet while a single fiber's bandwidth has increased by a factor of 10^6 since 1975, doubling every 16 months.

[0004] New processors and storage devices continue to break records in speed and capacity, supporting more powerful end devices throughout the network. However, computation continues to accumulate around data centers, which have to increase their workloads at a crippling pace, thus putting immense pressure on space and power consumption.

[0005] Finally, computer users in general are accustomed to computer systems that are deterministic and synchronous in nature, and think of such a structure as the norm. For example, when a browser issues a URL request for a Web page, the output is typically expected to appear shortly afterwards. It is also typically expected that everyone around the world will be able to retrieve the same page from the same Web server using the same URL.

[0006] The term peer-to-peer networking or computing (often referred to as P2P) may be applied to a wide range of technologies that greatly increase the utilization of information, bandwidth, and computing resources in the Internet. Frequently, these P2P technologies adopt a network-based computing style that neither excludes nor inherently depends on centralized control points. Apart from improving the performance of information discovery, content delivery, and information processing, such a style also can enhance the overall reliability and fault-tolerance of computing systems.

[0007] Peer-to-peer (P2P) computing, embodied by applications like Napster, Gnutella, and Freenet, has offered a compelling and intuitive way for Internet users to find and share resources directly with each other, often without requiring a central authority or server. As much as these diverse applications have broken new ground, they typically address only a single function, run primarily only on a single platform, and are unable to directly share data with other, similar applications.

[0008] Many peer-to-peer systems are built for delivering a single type of service. For example, Napster provides music file sharing, Gnutella provides generic file sharing, and AIM provides instant messaging. Given the diverse characteristics of these services and the lack of a common underlying P2P infrastructure, each P2P software vendor tends to create incompatible systems — none of them able to interoperate with one another. This means each vendor creates its own P2P user community, duplicating efforts in creating software and system primitives commonly used by all P2P systems. Moreover, for a peer to participate in multiple communities organized by different P2P implementations, the peer must support multiple implementations, each for a distinct P2P system or community, and serve as the aggregation point.

[0009] Many P2P systems today offer their features or services through a set of APIs that are delivered on a particular operating system using a specific networking protocol. For example, one system might offer a set of C++ APIs, with the system initially running only on Windows, over TCP/IP, while another system offers a combination of C and Java APIs, running on a variety of UNIX systems, over TCP/IP but also requiring HTTP. A P2P developer is then forced to choose which set of APIs to program to, and consequently, which set of P2P customers to target. Because there is little hope that the two systems will interoperate, if the developer wants to offer the same service to both communities, they have to develop the same service twice for two P2P platforms or develop a bridge system between them. Both approaches are inefficient and impractical considering the dozens of P2P platforms in existence.

[0010] Many P2P systems, especially those being offered by upstart companies, tend to choose one operating system

as their target deployment platform. The cited reason for this choice is to target the largest installed base and the fastest path to profit. The inevitable result is that many dependencies on platform-specific features are designed into (or just creep into) the system. This is often not the consequence of technical desire but of engineering reality with its tight schedules and limited resources.

[0011] This approach is clearly shortsighted. Even though the earliest demonstration of P2P capabilities are on platforms in the middle of the computing hardware spectrum, it is very likely that the greatest proliferation of P2P technology will occur at the two ends of the spectrum — large systems in the enterprise and consumer-oriented small systems. In fact, betting on any particular segment of the hardware or software system is not future proof.

[0012] Figures 1A and 1B are examples illustrating the peer-to-peer model. Figure 1A shows two peer devices 104A and 104B that are currently connected. Either of the two peer devices 104 may serve as a client of or a server to the other device. Figure 1B shows several peer devices 104 connected over the network 106 in a peer group. In the peer group, any of the peer devices 104 may serve as a client of or a server to any of the other devices.

[0013] Prior art peer-to-peer systems are generally built for delivering a single type of service, for example a music file sharing service, a generic file sharing service, or an instant messaging service. Given the diverse characteristics of these services and given the lack of a common underlying peer-to-peer infrastructure, each vendor tends to form various peer-to-peer "silos". In other words, the prior art peer-to-peer systems typically do not interoperate with each other. This means each vendor has to create its own peer-to-peer user community, duplicating efforts in creating primitives commonly used by peer-to-peer systems such as peer discovery and peer communication.

[0014] Discovery in a peer-to-peer environment may be based on centralized discovery with a centralized index. This method is used by such peer-to-peer applications as Napster and AIM. Discovery based on a centralized index may be efficient, deterministic, and well suited for a static environment. Such a method of discovery may also provide centralized control, provide a central point of failure, and provide easy denial of services. However, such a method of discovery may be expensive to scale and may degrade with aging.

[0015] Discovery in a peer-to-peer environment may also be based on net crawling. This method is used by such peer-to-peer applications as Gnutella and FreeNet. Discovery based on net crawling may be simple, adaptive, deterministic, inexpensive to scale, well suited for a dynamic environment, and may be difficult to attack. Such a method of discovery may also improve with aging. However, such a method of discovery may provide slower discovery than centralized control.

[0016] In a peer-to-peer environment, assume there is a peer-to-peer community offering a search capability for its members, where one member can post a query and other members can hear and respond to the query. One member is a Napster user and has implemented a feature so that, whenever a query is received seeking an MP3 file, this member will look up the Napster directory and then respond to the query with information returned by the Napster system. Here, a member without any knowledge of Napster may benefit because another member implemented a bridge to connect their peer-to-peer system to Napster. This type of bridging is very useful, but when the number of services is large, pair-wise bridging becomes more difficult and undesirable. Thus, it may be desirable to provide a platform bridge that may be used to connect various peer-to-peer systems together.

[0017] In another example, one engineering group requires a sizable storage capability, but also with redundancy to protect data from sudden loss. A common solution is to purchase a storage system with a large capacity and mirrored disks. Another engineering group later decides to purchase the same system. Both groups end up with a lot of extra capacity, and have to pay higher prices for the mirroring feature. Thus, it may be desirable to provide a mechanism by which each group may buy a simple storage system without the mirroring feature, where the disks can then discover each other automatically, form a storage peer group, and offer mirroring facilities using their spare capacity.

[0018] As yet another example, many devices such as cell phones, pagers, wireless email devices, Personal Digital Assistants (PDAs), and Personal Computers (PCs) may carry directory and calendar information. Currently, synchronization among the directory and calendar information on these devices is very tedious, if not impossible. Often, a PC becomes the central synchronization point, where every other device has to figure out a way to connect to the PC (using serial port, parallel port, IRDA, or other method) and the PC must have the device driver for every device that wishes to connect. Thus, it may be desirable to provide a mechanism by which these devices may interact with each other, without extra networking interfaces except those needed by the devices themselves, utilizing a common layer of communication and data exchange.

SUMMARY OF THE INVENTION

[0019] The accompanying claims relate to various aspects of the present invention.

[0020] Embodiments of a peer-to-peer platform for peer-to-peer computing on a network are described. In one embodiment, the peer-to-peer platform architecture may include, but is not limited to, protocols, advertisements, and core services. The peer-to-peer platform may include one or more protocols. To underpin this set of protocols, the peer-to-peer platform may define a number of concepts including peers, peer groups, advertisements, services, content, mes-

sages, pipes and pipe endpoints. Using the peer-to-peer platform protocols, peers may discover each other, communicate with each other, and cooperate with each other to form peer groups. In one embodiment, peer groups implicitly define a region scope that may limit peer propagation requests. Conceptually, a peer group may be viewed as a virtual entity that speaks the set of peer group protocols.

5 **[0021]** The peer-to-peer platform preferably provides a decentralized environment that minimizes single points of failure and is not dependent on any centralized services. Each protocol may be adopted into peer-to-peer services and applications. The peer-to-peer platform may be used to build a wide range of distributed services and applications in which every device is addressable as a peer, and where peers can bridge from one domain into another. The peer-to-peer platform may support distributed computing software that is flexible, interoperable, and available on any peer
10 on the network. The peer-to-peer platform may also support interoperable services and content on the Internet. The peers in a peer group may cooperate to provide a common set of services. Both centralized and decentralized services may be developed on top of the peer-to-peer platform. In the environment, services may be implemented to interoperate with other services giving rise to new P2P applications.

[0022] A system built on top of the peer-to-peer platform preferably functions in the same or similar fashion when
15 the system is expanded to a new networking environment or to a new class of devices, as long as there is a correct transport protocol handler for the new networking protocol. Thus, the peer-to-peer platform is preferably transport protocol independent. The protocols may be realized over networks including, but not limited to, the Internet, a corporate intranet, a dynamic proximity network, a home networking environment, LANs, and WANs. The protocols may also be realized within a single computer. The peer-to-peer platform is preferably independent of transport protocols. For example, the peer-to-peer platform may be implemented on top of TCP/IP, HTTP, Bluetooth, HomePNA, and other protocols.
20 The size and complexity of the network peers that may support these protocols preferably includes a wide range of peer implementations including peers implemented on, but not limited to, simple light switches, PDAs, cell phones, pagers, laptop and notebook computers, smart appliances, personal computers, workstations, complex, highly-available servers, mainframe computers and even supercomputers.

[0023] In one embodiment, the peer-to-peer platform may run on any of various operating systems including embedded operating systems (with the appropriate level of Java runtime support, if required) such as Windows95, 98,
25 2000, ME, and NT, Solaris, Unix, Macintosh, Linux, Java 2 Platform, Micro Edition (J2ME) and PersonalJava Technology. The peer-to-peer platform may be implemented in any of a variety of development environments using any of a variety of programming languages, or combinations of programming languages, including, but not limited to, Java, Java 2ME, C/C++, Perl, Python and KVM.
30

[0024] The peer-to-peer platform may include several layers. In one embodiment, the layers may include a core layer, a service layer and an application layer. In one embodiment, the peer-to-peer platform may include a core layer that defines and encapsulates minimal primitives that are common to peer-to-peer networking, including, but not limited to, peers, peer groups, peer discovery, peer communication, peer monitoring, and associated security primitives. This
35 layer may be shared by all peer-to-peer devices so that interoperability becomes possible.

[0025] The core layer provides core support for peer-to-peer services and applications. In a multi-platform, secure execution environment, the core mechanisms of peer groups, peer pipes and peer monitoring may be provided. Peer groups may establish a set of peers and naming within a peer group with mechanisms to create policies for creation and deletion, membership, advertising and discovery of other peer groups and peer nodes, communication, security,
40 and content sharing.

[0026] At the highest abstraction level, the peer-to-peer platform may be viewed as a set of protocols provided at the core layer. The peer-to-peer platform protocols may be used to provide and support ad hoc, pervasive, and multi-hop peer-to-peer (P2P) network computing. Using the protocols, peers can cooperate to form self-organized and self-configured peer groups independently of their positions in the network (e.g. edges, firewalls), and without the need of
45 a centralized management infrastructure. The peer-to-peer platform protocols may have very low overhead, make few assumptions about the underlying network transport and limited requirements of the peer environment, and may be used to deploy a wide variety of P2P applications and services in a highly unreliable and changing network environment. In one embodiment, the peer-to-peer platform may include core protocols including, but not limited to, a peer membership protocol, a peer discovery protocol, a peer resolver protocol, a peer information protocol, a pipe binding protocol, and a peer endpoint protocol. In one embodiment, peers are not required to implement all the core protocols. Further-
50 more, some peers may only use a portion (client-side or server-side only) of a protocol.

[0027] The peer-to-peer platform protocols may be implemented using a common messaging layer. This messaging layer binds the protocols to various network transports. Each protocol may be defined by one or more messages exchanged among participants of the protocol. Each message may have a predefined format, and may include various data fields. In one embodiment, the peer-to-peer platform may use XML as the encoding format. XML may provide
55 convenience in parsing and extensibility. Other embodiments of the peer-to-peer platform may use other encoding formats. The use of XML does not imply that all peer-to-peer platform peer nodes must be able to parse and to create XML documents.

[0028] The peer-to-peer platform may further include a peer-to-peer services layer. This layer may provide capabilities that may not be absolutely necessary for a peer-to-peer network to operate but that may be desirable to provided added functionality beyond the core layer in the peer-to-peer environment. The service layer may deal with higher-level concepts such as search and indexing, directory, storage systems, file sharing, distributed file systems, resource aggregation and renting, protocol translation, authentication and PKI (public key infrastructure) systems. These services, which may make use of the protocols and building blocks provided by the core layer, may be useful by themselves but also may be included as components in an overall P2P system. Thus, services may include one or more services provided by the peer-to-peer platform. These platform-provided services may include indexing, searching and file sharing services, for example. The services layer may provide hooks for supporting generic services (such as searching, sharing and added security) that may be used in many P2P applications. Thus, services may also include one or more services not provided as part of the peer-to-peer platform but rather provided by the peer-to-peer platform community.

[0029] The peer-to-peer platform may also include a peer-to-peer application layer. The application layer may support the implementation of integrated applications such as file sharing, resource sharing, monetary systems, distributed storage, peer-to-peer instant messaging, entertainment, content management and delivery, peer-to-peer email systems, distributed auction systems, among others. Applications may be "vertical" or they may be developed to interoperate with other distributed applications. One or more applications may be provided as part of the peer-to-peer platform. Applications may also include user-defined community applications not provided by the peer-to-peer platform. Applications may be built using peer services as well as the core layer. The peer-to-peer platform may support the fundamental levels broadly, and rely on the P2P development community to provide additional peer services and applications.

[0030] In one embodiment, the peer-to-peer protocols may use advertisements to describe and publish the existence of peer resources. An advertisement may be defined as a structured, language neutral metadata structure that names, describes, and publishes the existence of a peer-to-peer platform resource, such as a peer, a peer group, a pipe, or a service. In one embodiment, advertisements may be used in the peer-to-peer platform as language-neutral metadata structures. The peer-to-peer platform preferably includes advertisement documents to represent all of the peer-to-peer platform resources managed by the core platform, such as peers, peer groups, pipes and services. The peer-to-peer platform may define core advertisement types including, but not limited to, one or more of peer advertisements, peer group advertisements, pipe advertisements, service advertisements, content advertisements, and endpoint advertisements. In one embodiment, an advertisement is a markup language structured document that names, describes, and publishes the existence of a peer-to-peer platform resource. In one embodiment, peer-to-peer platform advertisements may be represented in the Extensible Markup Language (XML).

[0031] A typical peer-to-peer platform network may provide an inherently nondeterministic topology/response structure. In a peer-to-peer platform network, a specific resource request may not return for minutes, hours, or even days; in fact, it may never return at all. In addition, people from different parts of the world requesting the same resource are likely to get different copies of the resource from completely different locations. Peers may obtain content from multiple servers, ideally reaching a nearby one that is up and running. The original source peer need not service every resource request; in fact, it does not even have to be up and running.

[0032] The peer-to-peer platform provides the ability to replicate information toward end users. Popular content tends to be replicated more often, making it easier to find as more copies are available. Peers do not have to always go back to the same peer to obtain the information they want, as is typical in the client/server model. Peers may obtain information from neighboring peers that have already cached the information. Each peer may become a provider to all other peers.

[0033] In one embodiment the peer-to-peer platform may enable peers to find content that is closest (local) to their network location. This content may include data (e.g. files) or even services and applications. For example, if a peer node in an office peer-to-peer network using the peer-to-peer platform is moved, the peer-to-peer platform may allow the peer to automatically locate content including services hosted by other peers closest to the peer's new location, without requiring any manual reconfiguration. Further, at least some content may be copied or moved to the peer in its new location and/or to other peers local to the new location.

[0034] In one embodiment, the peer-to-peer platform may support and/or provide sophisticated naming and binding services. In one embodiment, the peer-to-peer platform may use a universal unique identifier (UUID) to refer to an entity (e.g. a peer, peer group, pipe, content, etc.). In one embodiment, the UUIDs may be used in providing flexible configuration and seamless relocation of peer nodes on a peer-to-peer network, and may assist in locating and accessing content including services nearest to a peer node when the peer node is moved. The peer-to-peer protocols and UUIDs may provide the ability for peer nodes to move to different peer groups and/or peer regions and access services and other content independent of network physical addresses and without requiring reconfiguration of the peer node. Peer nodes may be relocated and access services and other content that are locally hosted or services and other content hosted in their original peer group if the services and other content are not required to be locally hosted.

BRIEF DESCRIPTION OF THE DRAWINGS

[0035] Particular embodiments of the invention are described hereinafter with reference to the accompanying drawings in which like reference signs refer to like elements:

Figure 1A illustrates a prior art example of two devices that are currently connected as peers;
 Figure 1B illustrates a prior art example of several peer devices connected over the network in a peer group;
 Figure 2 illustrates one embodiment of peer-to-peer platform software architecture at the conceptual level;
 Figure 3 illustrates an exemplary content identifier according to one embodiment;
 Figure 4 illustrates a point-to-point pipe connection between peers according to one embodiment;
 Figure 5 illustrates a peer-to-peer platform message format according to one embodiment;
 Figure 6 illustrates the content of a peer advertisement according to one embodiment;
 Figure 7 illustrates the content of a peer group advertisement according to one embodiment;
 Figure 8 illustrates the content of a pipe advertisement according to one embodiment;
 Figure 9 illustrates the content of a service advertisement according to one embodiment;
 Figure 10 illustrates the content of a content advertisement according to one embodiment;
 Figure 11 illustrates the content of an endpoint advertisement according to one embodiment;
 Figure 12 illustrates protocols and bindings in a peer-to-peer platform according to one embodiment;
 Figure 13 illustrates discovery through a rendezvous proxy according to one embodiment;
 Figure 14 illustrates discovery through propagate proxies according to one embodiment;
 Figure 15 illustrates using messages to discover advertisements according to one embodiment;
 Figure 16 illustrates one embodiment of using peer resolver protocol messages between a requesting peer and a responding peer;
 Figure 17 illustrates one embodiment of using peer information protocol messages between a requesting peer and a responding peer;
 Figure 18 illustrates several core components and how they interact for discovery and routing according to one embodiment;
 Figure 19 illustrates one embodiment of message routing in a peer-to-peer network that uses the peer-to-peer platform;
 Figure 20 illustrates traversing a firewall in a virtual private network when access is initiated from outside only according to one embodiment;
 Figure 21 illustrates email exchange through an email gateway according to one embodiment;
 Figure 22 illustrates traversing a firewall when access is initiated from the inside according to one embodiment;
 Figure 23 illustrates embodiments of a peer-to-peer platform proxy service, and shows various aspects of the operation of the proxy service;
 Figure 24 illustrates a method of using a proxy service for peer group registration according to one embodiment;
 Figure 25 illustrates peer group registration across a firewall according to one embodiment;
 Figure 26 illustrates a method of providing peer group membership through a proxy service according to one embodiment;
 Figures 27A and 27B illustrate a method of providing privacy in the peer-to-peer platform according to one embodiment;
 Figures 28A and 28B illustrate one embodiment of a method for using a peer-to-peer platform proxy service as a certificate authority;
 Figure 29 is a flowchart illustrating a peer accessing content from different locations according to one embodiment;
 Figure 30 is a flowchart illustrating a peer accessing content from different locations according to one embodiment.

[0036] While the invention is described herein by way of example for several embodiments and illustrative drawings, those skilled in the art will recognize that the invention is not limited to the embodiments or drawings described. It should be understood, that the drawings and detailed description thereto are not intended to limit the invention to the particular form disclosed, but on the contrary, the intention is to cover all modifications, equivalents and alternatives falling within the scope of the present invention. The headings used herein are for organizational purposes only and are not meant to be used to limit the scope of the description or the claims. As used throughout this application, the word "may" is used in a permissive sense (i.e., meaning having the potential to), rather than the mandatory sense (i.e., meaning must). Similarly, the words "include", "including", and "includes" mean including, but not limited to.

DETAILED DESCRIPTION OF EMBODIMENTS OF THE INVENTION

[0037] Embodiments of an open network computing platform designed for peer-to-peer computing are described. The network computing platform may be referred to as a peer-to-peer platform. The peer-to-peer platform may be used to build a wide range of distributed services and applications in which every device is addressable as a peer, and where peers can bridge from one domain into another. The peer-to-peer platform may enable developers to focus on their own application development while easily creating distributed computing software that is flexible, interoperable, and available on any peer on the expanded Web. The peer-to-peer platform may enable software developers to deploy interoperable services and content, further spring-boarding the peer-to-peer revolution on the Internet. The peer-to-peer platform addresses the problems of prior art peer-to-peer systems by providing a generic and service-agnostic peer-to-peer platform that may be preferably defined by a small number of protocols. Each protocol is preferably easy to implement and easy to be adopted into peer-to-peer services and applications. Thus, service offerings from one vendor may be used, perhaps transparently, by the user community of another vendor's system.

[0038] The peer-to-peer platform extends P2P computing to enable a wide range of distributed computing applications and overcome the limitations typically found in prior art P2P applications. The peer-to-peer platform is a network computing technology that provides a set of simple, small, and flexible mechanisms that can support P2P computing on any platform, anywhere, and at any time. The peer-to-peer platform generalizes P2P functionality and provides core technology that addresses the limitations of prior art P2P computing technologies.

[0039] The peer-to-peer platform is a modular platform that provides simple and essential building blocks for developing a wide range of distributed services and applications. The peer-to-peer platform specifies a set of protocols rather than an API. Thus, the peer-to-peer platform can be implemented in any language on any Operating System to provide solutions ranging from providing a simple protocol-based wrapper that enables a small device to join a network of peers to developing a fully integrated application that supports metering, monitoring, high-level security and communication across server-class systems.

[0040] In one embodiment, the peer-to-peer platform architecture may include, but is not limited to, protocols, advertisements, and core services. Network protocol bindings may be used to preferably ensure interoperability with existing content transfer protocols, network transports, routers, and firewalls. The peer-to-peer platform may be used to combine network nodes (peers) into a simple and coherent peer-to-peer network computing platform. The platform is preferably directed at providing several benefits including, but not limited to, no single point of failure, asynchronous messaging, the ability for peers to adapt to their network environment, and moving content towards its consumers.

[0041] Figure 2 illustrates one embodiment of peer-to-peer platform software architecture at the conceptual level. The peer-to-peer platform may include several layers. In one embodiment, the software stack may be described using three layers; a peer-to-peer platform (core) layer 120, a service layer 140 and an application layer 150. In one embodiment, the peer-to-peer platform may include a core layer 120 that defines and encapsulates minimal primitives that are common to peer-to-peer networking, including, but not limited to, peers 110, peer groups 122, peer discovery 124, peer communication (e.g. pipes) 126, peer monitoring 128, and associated security primitives 130. This layer may be shared by all peer-to-peer devices so that interoperability becomes possible.

[0042] A peer may be defined as any entity that runs some or all of one or more protocols provided by the peer-to-peer platform core layer. As such, a peer may manifest in the form of a processor, a process or a device. A peer may be anything with a digital heartbeat that supports the peer-to-peer platform core, including sensors, servers, PCs, computers up to and including supercomputers, PDAs, manufacturing and medical equipment, phones and cellular phones. In order to interact with other peers (e.g. to form or join peer groups), the peer needs to be connected to some kind of network (wired or wireless), such as IP, Bluetooth, or Havi, among others.

[0043] The peer-to-peer platform may provide mechanisms through which peers may discover each other, communicate with each other, and cooperate with each other to form peer groups. Peers may discover each other on the network to form transient or persistent relationships called peer groups. A peer group is a collection of peers connected by a network that share a common set of interests and that have agreed upon a common set of rules to publish, share and access any computer content (code, data, applications, or other collections of computer representable resources), and communicate among themselves. Peer groups may also be statically predefined. The peers in a peer group may cooperate to provide a common set of services. A peer group may be viewed as an abstract region of the network, and may act as a virtual subnet. The concept of a region virtualizes the notion of routers and firewalls, subdividing the network in a self-organizing fashion without respect to actual physical network boundaries. In one embodiment, peer groups implicitly define a region scope that may limit peer propagation requests. Conceptually, a peer group may be viewed as a virtual entity that speaks the set of peer group protocols.

[0044] A peer group may theoretically be as large as the entire connected universe. Naming anything uniquely is a challenge in such a large namespace. In one embodiment, the peer-to-peer platform may support and/or provide sophisticated naming and binding services. In one embodiment, the peer-to-peer platform may use a universal unique identifier (UUID), for example, a 64- or 128-bit datum, to refer to an entity (e.g. a peer, peer group, pipe, content, etc.).

For example, UUIDs may be embedded in advertisements for internal use. UUIDs preferably may be used to guarantee that each entity has a unique UUID within a local runtime environment and serves as a canonical way of referring to an entity, but because a global state is not assumed, it may not be possible to provide a guarantee of uniqueness across an entire community that may consist of millions of peers. This may not be a problem because a UUID may be used within the peer-to-peer platform as an internal identifier. This may become significant only after the UUID is securely bound to other information such as a name and a network address. In one embodiment, Uniform Resource Name (URN) format may be used for the expression of UUIDs.

[0045] The core layer 120 provides core support for peer-to-peer services and applications. In a multi-platform, secure execution environment, the core mechanisms of peer groups, peer pipes and peer monitoring may be provided. Peer groups 122 may establish a set of peers and naming within a peer group with mechanisms to create policies for creation and deletion, membership, advertising and discovery of other peer groups and peer nodes, communication, security, and content sharing. Pipes provide virtual communication channels among peers. Messages sent in pipes may support transfer of data, content, and code in a protocol-independent manner, allowing a range of security, integrity, and privacy options. In one embodiment, messages may be structured with a markup language such as XML. Peer monitoring 128 enables control of the behavior and activity of peers in a peer group and can be used to implement peer management functions including access control, priority setting, traffic metering, and bandwidth balancing.

[0046] The core layer 120 may include protocols and building blocks to enable key mechanisms for peer to peer networking, including discovery, transport (including firewall handling and limited security), and the creation of peers and peer groups. The core layer 120 is preferably thin and small, and preferably provides interesting and powerful primitives for use by services and applications in the other layers. The core layer 120 may support choices such as anonymous vs. registered users and encrypted vs. clear text content without imposing specific policies on developers. Policy choices may be made, or when necessary, implemented, at the service layer 140 and/or application layer 150. For example, administration services such as accepting or rejecting a peer's membership in a peer group may be implemented using the functionality provided by the core layer 120.

[0047] The core components of the peer-to-peer protocol may be used to implement discovery mechanisms for searching, publishing and recovering of core abstractions (e.g. peers, peer group, pipes, endpoints, and advertisements). These mechanisms are preferably simple, administration free, and do not require special peers to act as "master" peers. These mechanisms may allow processes in the peer-to-peer network, in absence of help from other applications and/or services, to bootstrap and find out the information necessary to access applications and services that can help. Also, the core may "return" to this standalone behavior and still function if helper applications or services fail. In one embodiment, safety mechanisms may be put in place in order to avoid a major overflow of "web-crawling". In one embodiment, applications and/or services that support the peer-to-peer protocol may access, control, and/or override the core components, even to the extreme of implementing a centralized, client-server model based on the core components.

[0048] At the highest abstraction level, the peer-to-peer platform may be viewed as a set of protocols provided at the core layer 120. In one embodiment, a common thread among peer-to-peer platform peers is protocols, not APIs or software implementations. The peer-to-peer platform protocols preferably guarantee interoperability between compliant software components executing on potentially heterogeneous peer runtimes. Thus the peer-to-peer platform is preferably agnostic to programming languages. The term compliant may refer to a single protocol only. That is some peers may not implement all the core protocols. Furthermore, some peers may only use a portion (client-side or server-side only) of a protocol.

[0049] Each protocol may be defined by one or more messages exchanged among participants of the protocol. Each message may have a predefined format, and may include various data fields. In one embodiment the protocols may utilize messaging such as XML messages. The peer-to-peer platform connects peer nodes with each other. The peer-to-peer platform is preferably platform-independent by virtue of being a set of protocols. As such, the peer-to-peer platform may not require APIs and remains independent of programming languages, so that it can be implemented in C/C++, Java, Java 2ME, Perl, Python or other languages. This means heterogeneous devices with completely different software stacks can preferably interoperate through the peer-to-peer platform protocols. To underpin this set of protocols, the peer-to-peer platform may define a number of concepts including peer, peer group, advertisement, message, pipe, and more.

[0050] In one embodiment, peer-to-peer protocols may be embodied as markup language (e.g. XML) messages that may be sent between two peers. In one embodiment, the peer-to-peer platform messages may define the protocols used to discover and connect peers and peer groups, and to access resources offered by peers and peer groups, among others. The use of markup language (e.g. XML) messages to define protocols may allow many different kinds of peers to participate in a protocol. Each peer may be free to implement the protocol in a manner best suited to its abilities and role. For example, not all peers are capable of supporting a Java runtime environment. In one embodiment, the protocol definition does not require nor imply the use of Java on a peer.

[0051] Several peer-to-peer platform protocols that may be provided by embodiments of the peer-to-peer platform

are described later in this document. The protocols defined in this document may be realized over networks including, but not limited to, the Internet, a corporate intranet, a dynamic proximity network, a home networking environment, LANs, and WANs. The protocols defined in this document may also be realized within a single computer. Thus, the peer-to-peer platform is preferably transport protocol independent. The size and complexity of the network peers that may support these protocols preferably includes a wide range of peer implementations including peers implemented on, but not limited to, simple light switches, PDAs, cell phones, pagers, laptop and notebook computers, smart appliances, personal computers, workstations, complex, highly-available servers, mainframe computers and even super-computers.

[0052] The peer-to-peer platform may further include a peer-to-peer services layer 140. This layer may provide capabilities that may not be absolutely necessary for a peer-to-peer network to operate but that may be desirable to provided added functionality beyond the core layer 120 in the peer-to-peer environment. The service layer 140 may deal with higher-level concepts such as search and indexing, directory, storage systems, file sharing, distributed file systems, resource aggregation and renting, protocol translation, authentication and PKI (public key infrastructure) systems. These services, which may make use of the protocols and building blocks provided by the core layer 120, may be useful by themselves but also may be included as components in an overall P2P system. Thus, services may include one or more services 144 provided by the peer-to-peer platform. These platform-provided services 144 may include indexing, searching and file sharing services, for example. The services layer 140 may provide hooks for supporting generic services (such as searching, sharing and added security) that are used in many P2P applications. Thus, services may also include one or more services 142 not provided as part of the peer-to-peer platform but rather provided by the peer-to-peer platform community. These services 142 may be user-defined and may be provided, for example, to member peers in a peer group as a peer group service.

[0053] Services may expand upon the capabilities of the core layer 120 and may be used to facilitate application development. Facilities provided as services in the service layer 140 may include mechanisms for search and indexing, directory, storage systems, file sharing, distributed file systems, resource aggregation and renting, protocol translation, authentication, PKI services, and caching code and content to enable cross-application bridging and translation of files, among others. Searching capabilities may include distributed, parallel searches across peer groups that are facilitated by matching an XML representation of a query to be processed with representations of the responses that can be provided by each peer. These facilities may be used for simple searches, for example searching a peer's repository, or more complex searches of dynamically generated content that is unreachable by conventional search engines. P2P searches may be conducted across a company's intranet, for example, to quickly locate relevant information within a secure environment. By exercising tight control over peer group membership and enabling encrypted communication between peers, a company may extend this capability to its extranet, including business partners, consultants, and suppliers as peers. The same mechanisms that facilitate searches across the peer group may be used as a bridge to incorporate Internet search results, and to include data outside of the peer's own repository, for example searching a peer's disk. The peer services layer 140 may be used to support other custom, application-specific functions. For example, a secure peer messaging system may be built to allow anonymous authorship and a persistent message store. The peer services layer 140 provides the mechanisms to create such secure tools; the application developers themselves may determine specific tool policies.

[0054] The peer-to-peer platform may also include a peer-to-peer application layer 150. The application layer 140 may support the implementation of integrated applications such as file sharing, resource sharing, monetary systems, distributed storage, peer-to-peer instant messaging, entertainment, content management and delivery, peer-to-peer email systems, distributed auction systems, among others. Applications may be "vertical" or they may be developed to interoperate with other distributed applications. One or more applications 154 may be provided as part of the peer-to-peer platform. For example, one embodiment of the peer-to-peer platform may include a shell application 160 as a development environment built on top of the platform. The shell application may provide interactive access to the peer-to-peer platform via a simple command line interface 162.

[0055] Applications may also include community applications 152 not provided by the peer-to-peer platform. These community applications 152 may be user-defined and may be provided, for example, to member peers in a peer group as a peer group application.

[0056] In one embodiment, the boundary between services and applications is not rigid. An application to one customer can be viewed as a service to another customer. An application may use services. Services may serve as protocols that may be shared among various applications. An application may provide a user interface, a way to define a set of files to share, a way to initiate a search, a way to display the results, and a way to initiate a file transfer, for example. Such an application may make use of a set of services, for example a reliable point-to-point file transfer service, a distributed search service, and a discovery service to locate other peers, among others.

[0057] Applications may be built using peer services as well as the core layer 120. The peer-to-peer platform may support the fundamental levels broadly, and rely on the P2P development community to provide additional peer services and applications. Peer applications enabled by both the core layer 120 and peer services layer 140 may include P2P

5 auctions that link buyers and sellers directly, with buyers able to program their bidding strategies using a simple scripting language, for example. Resource-sharing applications, such as SETI@home, may be built more quickly and easily, with heterogeneous, worldwide peer groups supported from day one. Instant messaging, mail, and calendaring services may facilitate communication and collaboration within peer groups that are secure and independent of service provider-hosted facilities. Virtually any other type of application may be build on top of the core layer 120 and services layer 140.

[0058] Some features, such as security, may manifest in all three layers and throughout a P2P system, albeit in different forms according to the location in the software architecture. The system is preferably modular, and allows developers to pick and choose a collection of services and applications that suits their needs.

[0059] A typical peer-to-peer platform network may provide an inherently nondeterministic topology/response structure. In a peer-to-peer platform network, a specific resource request may not return for minutes, hours, or even days; in fact, it may never return at all. In addition, people from different parts of the world requesting the same resource are likely to get different copies of the resource from completely different locations. Peers may obtain content from multiple servers, ideally reaching a nearby one that is up and running. The original source peer need not service every resource request; in fact, it does not even have to be up and running. The nondeterministic structure may also help provide the optimized use of network bandwidth. The concentrated localized traffic congestion typical of today's Web doesn't affect P2P networking. The nondeterministic structure may also help provide a lowered cost of content distribution. The P2P network can absorb contents and replicate it for easy access. The nondeterministic structure may also help provide leveraged computing power from every node in the network. With asynchronous operations, a user may issue many requests for many resources or services simultaneously and have the network do the work. The nondeterministic structure may also help provide unlimited scalability. A properly designed P2P application may span the entire known connected universe without hitting scalability limits; this is typically not possible with centralized schemes. Note, however, that the peer-to-peer platform also may support deterministic, synchronous applications.

[0060] As an example of a nondeterministic, asynchronous application, consider a network-based music request service that operates over a peer-to-peer platform-based P2P network. A peer submits multiple requests for music files and then checks back later to see if the music request service in the peer group has found them. A few requested files have been found, but others cannot be located. The service's response in regards to the files that cannot be located may be something like "Music selection and availability changes continuously; please retry your request later." This is an acceptable nondeterministic outcome. Even though the service couldn't find a file, the same file may be available later if the same request is resubmitted, because peers that host the desired files may have come online in the meantime.

[0061] The peer-to-peer platform provides the ability to replicate information toward end users. Popular content tends to be replicated more often, making it easier to find as more copies are available. Peers do not have to always go back to the same peer to obtain the information they want, as is typical in the client/server model. Peers may obtain information from neighboring peers that have already cached the information. Each peer may become a provider to all other peers.

[0062] In one embodiment the peer-to-peer platform may enable peers to find content that is closest to them. This content may include data (e.g. files) or even services and applications. For example, if a peer node in an office peer-to-peer network using the peer-to-peer platform is moved, the peer-to-peer platform may allow the peer to automatically locate content (e.g. using a discovery service that participates in the discovery protocol) including services (e.g. a printer service and an email service) hosted by other peers closest to the peer's new location, without requiring any manual reconfiguration. Further, at least some content may be copied or moved to the peer in its new location and/or to other peers proximate to the new location.

[0063] Figure 29 is a flowchart illustrating a peer accessing content from different locations according to one embodiment. A peer that participates in the peer-to-peer platform may couple to a network at a network location through a network interface as indicated at 400. On the network, a plurality of peers may include an instance of a content that the peer may require. The peer may access instances of the content local to the network location of the peer on the network as indicated at 402. Localness may be determined, in one embodiment, by the number of jumps on the network required to access the content. In one embodiment, member peers in a peer group that the peer is a member of may be "nearest" peers. The peer may use a discovery method as described herein to discover the instances of the content local to the peer. As indicated at 404, the peer may uncouple from the network at the network location and move to a different network location and recouple to the network. The different network location may be at a different physical location or alternatively a different address at the same physical location. The peer may then discover and access a different instance of the content provided by a different peer (or the same peer) local to the different network location as indicated at 406.

[0064] In one embodiment, the UUIDs may be used in providing flexible configuration and seamless relocation of peer nodes on a peer-to-peer network, and may assist in locating and accessing content including services nearest to a peer node when the peer node is moved. For example, a businessperson based in New York may participate in a peer-to-peer network based on the peer-to-peer protocols using a notebook computer or other portable computing

device connected to a LAN as a peer node. The businessperson may access an instance of an email and/or other services locally hosted by other peer nodes in a peer group on the LAN. If the businessperson travels to Paris, for example, and takes the notebook computer, the notebook computer may be connected to a different LAN at the Paris location and participate in the peer-to-peer network. Because the peer node has a unique ID in the peer-to-peer network (the UUID) rather than just a static network address, the peer node may seamlessly access instances of an email service and other services locally hosted on the LAN, or alternatively hosted on a peer node at the peer node's original location or elsewhere, using the UUID to establish its identity. The peer node may rejoin the peer group in New York to access one or more instances of services and other content hosted on the peer group, and may also join a peer group at the Paris location to access one or more other instances of services and content. Thus, the peer-to-peer protocols and UUIDs may provide the ability for peer nodes to move to different peer groups and/or peer regions and access services and other content independent of network addresses and without requiring reconfiguration of the peer node. For example, when the exemplary peer node moves to Paris, connects to the network (at a different network address) and accesses an instance of an email service (either locally or remotely hosted, for example in the New York peer group), the email service may identify the peer node by its unique ID and route the peer's email to the peer node at the new network address without requiring reconfiguration of the peer node. Thus, peer nodes may be relocated and access services and other content that are locally hosted or services and other content hosted in their original peer group if the services and other content are not required to be locally hosted.

[0065] Figure 30 is a flowchart illustrating a peer accessing content from different network locations according to one embodiment. A peer that participates in the peer-to-peer platform may couple to a network at a network location through a network interface as indicated at 410. On the network, a plurality of peers may include an instance of a service that the peer may require. The peer may access instances of the service local to the network location of the peer on the network as indicated at 412. Localness may be determined, in one embodiment, by the number of jumps on the network required to access the content. In one embodiment, member peers in a peer group that the peer is a member of may be "local" peers. The peer may use a discovery method as described herein to discover the instances of the service local to the peer. As indicated at 414, the peer may uncouple from the network at the network location and move to a different network location and recouple to the network. The different network location may be at a different physical location or alternatively a different address at the same physical location. The peer may then access an instance of the service from the new network location as indicated at 416. The peer may discover and access a different instance of the service provided by a different peer (or the same peer) local to the different network location. Alternatively, the peer may access the instance of the service local to the old network location, even though the instance of the service is not the most local instance of the service. In one embodiment, the peer may provide a unique identifier to the accessed instance of the service to distinguish the peer node from the other peer nodes on the network. The different instance of the service may then route information (e.g. messages, email, etc.) to the peer at the different network location. Thus, a peer may be moved on a network to different network locations or addresses, and the unique identifier may be provided to instances of services to seamlessly route information to the peer node.

[0066] Note that the two methods described in Figures 29 and 30 are merely examples. The methods may be implemented in software, hardware, or a combination thereof. The order of method may be changed, and various elements may be added, reordered, combined, omitted, modified, etc.

[0067] The peer-to-peer platform preferably provides a decentralized environment that minimizes single points of failure and is not dependent on any centralized services. Both centralized and decentralized services may be developed on top of the peer-to-peer platform. With the addition of each new network peer, the network platform preferably becomes more robust as it expands. In the environment, services may be implemented to interoperate with other services giving rise to new P2P applications. For example, a P2P communications service like instant messaging may easily be added to a resource-sharing P2P application if both support at least the necessary peer-to-peer platform protocols.

[0068] The peer-to-peer platform may provide interoperability. The peer-to-peer platform may be used by developers independent of preferred programming languages, development environments, or deployment platforms. Embodiments of the peer-to-peer platform may enable interconnected peers to easily locate each other, communicate with each other, participate in community-based activities, and offer services to each other seamlessly across different P2P systems and different communities. The peer-to-peer platform may also provide platform independence. Embodiments of the peer-to-peer platform may be independent of programming languages (such as C/C++, Java, Perl, and KVM), system platforms (such as the Microsoft Windows, UNIX®, Solaris, Linux and Macintosh platforms), and networking platforms (such as TCP/IP, Bluetooth and Havi). Thus, heterogeneous devices with completely different software stacks may interoperate through the peer-to-peer platform protocols. Embodiments of the peer-to-peer platform may be implementable on any device with a digital heartbeat, including, but not limited to, sensors, consumer electronics, Personal Digital Assistants (PDAs), appliances, network routers, desktop computers, data-center servers, and storage systems. Embodiments of the peer-to-peer platform may enable peers, independent of software and hardware platform, to benefit and profit from being connected to millions of other peers.

[0069] In one embodiment, the peer-to-peer platform may run on any of various operating systems including em-

bedded operating systems (with the appropriate level of Java runtime support, if required) such as Windows95, 98, 2000, ME, and NT, Solaris, Unix, Macintosh, Linux, Java 2 Platform, Micro Edition (J2ME) and Personal Java Technology. The peer-to-peer platform may be implemented in any of a variety of development environments using any of a variety of programming languages, or combinations of programming languages, including, but not limited to, Java, Java 2ME, C/C++, Perl, Python and KVM. In one embodiment, the peer-to-peer platform may be implemented in Java. In one embodiment, a peer-to-peer platform may be implemented in C/C++ on some devices, for example, to support devices without Java support. In one embodiment, a peer-to-peer platform may be implemented in KVM on some devices, for example, so that all KVM capable devices such as PDAs and cell phones can be peer-to-peer platform peers. Programming languages other than those listed may also be used in various embodiments.

[0070] A minimal device with the ability to generate a text string may theoretically participate in a peer-to-peer platform network (though not necessarily in every P2P application). The simplistic device may need a surrogate peer on the P2P network. This surrogate peer may perform discovery, advertisement, and communications on behalf of the simplistic device (or many simplistic devices). The location of the surrogate may be hard-wired into the simplistic device. In this way, the simplistic device with the help of the surrogate can be a full-fledged peer on the peer-to-peer platform network. For example, a GPS locator, strapped to a sea turtle and sending out peer-to-peer platform messages wirelessly with location information, may become a peer on a peer-to-peer platform network.

[0071] The peer-to-peer platform is preferably independent of transport protocols. For example, the peer-to-peer platform may be implemented on top of TCP/IP, HTTP, Bluetooth, HomePNA, and other protocols. Thus, a system built on top of the peer-to-peer platform preferably functions in the same or similar fashion when the system is expanded to a new networking environment or to a new class of devices, as long as there is a correct transport protocol handler for the new networking protocol.

[0072] In one embodiment, the peer-to-peer platform may use XML as the encoding format. XML may provide convenience in parsing and extensibility. Other embodiments of the peer-to-peer platform may use other encoding formats. The use of XML does not imply that all peer-to-peer platform peer nodes must be able to parse and to create XML documents. For example, a cell phone with limited resources may be programmed to recognize and to create certain canned XML messages and can still participate in a peer-to-peer platform network of peers. In one embodiment, a lightweight XML parser may be used that supports a subset of XML. This may help reduce the size of the peer-to-peer platform.

[0073] There may be areas in a peer-to-peer environment where there is not one correct way to do something or where what should be done depends on the nature and context of the overriding application. For example, in the area of security, every P2P application may choose a different authentication scheme, a different way to ensure communication security, a different encryption algorithm for data security, a different signature scheme for authenticity, and a different access control policy. Therefore, for these areas, the peer-to-peer platform may preferably focus on mechanisms instead of policy, so that application developers can have the maximum freedom to innovate and offer competitive solutions.

[0074] Implementations of the peer-to-peer platform may be illustrated with a few application or usage scenarios. For example, assume there is a peer-to-peer community offering a search capability for its members, where one member can post a query and other members can hear and respond to the query. One member is a Napster user and has implemented a feature so that, whenever a query is received seeking an MP3 file, this member will look up the Napster directory and then respond to the query with information returned by the Napster system. Here, a member without any knowledge of Napster may benefit because another member implemented a bridge to connect their peer-to-peer system to Napster. The peer-to-peer platform may provide a platform bridge that may be used to connect the various peer-to-peer systems together.

[0075] In another example, one engineering group requires a sizable storage capability, but also with redundancy to protect data from sudden loss. Using the peer-to-peer platform, each group may buy a simple storage system without a mirroring feature, where the disks can then discover each other automatically, form a storage peer group, and offer mirroring facilities using their spare capacity.

[0076] As yet another example, many devices such as cell phones, pagers, wireless email devices, Personal Digital Assistants (PDAs), and Personal Computers (PCs) may carry directory and calendar information. Using the peer-to-peer platform, these devices may be able to interact with each other, without extra networking interfaces except those needed by the devices themselves, using the peer-to-peer platform as the common layer of communication and data exchange.

Peers

[0077] Network nodes (peers) of various kinds may join the peer-to-peer networking platform by implementing one or more of the platform's protocols. Each peer operates independently and asynchronously of any other peer, providing a degree of reliability and scalability not typically found in current distributed systems. Some peers may have more

dependencies with other peers due to special relationships (e.g. gateways or routers). In one embodiment, a peer does not need to understand all of the protocols of the peer-to-peer platform. The peer can still perform at a reduced level if it does not support one or more of the protocols.

[0078] Peers may publish and provide network resources (e.g. CPU, storage and routing resources) that may be used by other peers. Peers typically interact with a small number of other peers (network neighbors or buddy peers). Peers that provide the same set of services tend to be inter-changeable. Thus, it may not matter which peers a peer interacts with. Generally, assumptions should not be made about peer reliability or connectivity, as a peer may appear or leave the network at any time. Peers may have persistent storage. A peer may optionally cache information.

[0079] Peers may have multiple network interfaces, though preferably a peer does not need to publish all of its interfaces for use with the peer-to-peer protocols. Each published interface may be advertised as a peer endpoint. In one embodiment, a peer endpoint is an identifier (e.g. a URN or URI) that uniquely identifies a peer network interface. Peer endpoints may be used by peers to establish direct point-to-point connection between two peers. Peers may not have direct point-to-point network connection between themselves, either due to lack of physical network connections, or network configuration (NATs, firewalls, proxies, etc.), and thus a peer may have to use one or more intermediary peers to route a message from an endpoint to another peer endpoint.

[0080] The term rendezvous peer may be used to designate a peer that is designated to be a rendezvous point for discovering information about other peers, peer groups, services and pipes. Rendezvous peers preferably cache information that may be useful to peers including new peers. Rendezvous peers may provide an efficient mechanism for peers that are far away to find (e.g. discover) each other. Rendezvous peers may make peer discovery more practical and efficient. Preferably, a peer group is not required to have a rendezvous peer. In one embodiment, any or even all members of a peer group may become rendezvous peers in a peer group. In one embodiment, each peer group may have different policies to authorize a peer to become a rendezvous peer.

[0081] The term router peer may be used to describe a peer that crosses one or more regions and that is designated to be a router between the regions. Router peers may be used to route messages between different network protocols (e.g. TCP/IP, Irda) or to peers that are behind firewalls. In one embodiment, any or all peer members may become routers. In one embodiment, peer groups may have different policies to authorize a peer to become a router peer for other peers.

[0082] Peers may be identified by their unique ID (UUID) rather than by a fixed address. When a peer boots, it attempts to contact other peers. In one embodiment, contacted peers may include variable-sized caches that map nearby peers' UUID to their current address. This allows embodiments of the peer-to-peer platform to be run over a dialup connection, for example.

[0083] In one embodiment, a peer may be assigned a unique string as a name. Any naming scheme may be used. In one embodiment, names are not unique unless a coordinated naming service is used to guarantee name uniqueness. A naming service is typically a centralized service that guarantees the uniqueness of name and can be used to register name mapping. Examples of naming services are DNS and LDAP. Use of a naming service is preferably optional.

Peer Groups

[0084] Preferably, the peer-to-peer platform describes how to create and discover peer groups, but does not dictate when, where, or why to create a peer group, the type of the group, or the membership of the group. A peer group may provide a common membership definition. Each peer group may establish its own membership policy in a range from open (any peer can join) up to highly secure and protected (a peer may join only if it possesses sufficient credentials).

[0085] In one embodiment, peers wishing to join a peer group may first locate a current member, and then request to join the peer group. The peer-to-peer platform may define how to discover peer groups, e.g. using a peer discovery protocol. The application to join may be rejected or accepted by the collective set of current members in accordance with the peer group's membership policy. In one embodiment, a peer group core membership service may be used to enforce a vote among one or more group members. Alternatively, one or more group representative member peers may be elected or appointed to accept or reject new membership applications.

[0086] In one embodiment, the peer-to-peer platform is not concerned with what sequence of events a peer or a peer group comes into existence. Moreover, in one embodiment, the peer-to-peer platform does not limit how many groups a peer can belong to. In one embodiment, nested and/or overlapping peer groups may be formed. In one embodiment, there may be a special group, called the World Peer Group, which may include all peer-to-peer platform peers. The world peer group preferably provides the minimum seed for every peer to potentially find each other and form new groups. In one embodiment, the world peer group has an open membership policy (e.g. has a null membership authenticator service). Some peers inside the world peer group may not be able to discover or communicate with each other — e.g., they may be separated by a network partition. In one embodiment, participation in the World Peer Group is by default.

[0087] The peer-to-peer platform may use the concept of a peer group as an implicit scope of all messages originated

from within the group. Peer groups may serve to subdivide the network into abstract regions providing an implicit scoping mechanism. Peer groups may provide a limited scoping environment to ensure scalability. Peer groups may be formed and self organized based upon the mutual interest of peers. In one embodiment, no particular rules are imposed on the way peer groups are formed, but peers with the same interests may tend to join the same peer groups.

5 **[0088]** In one embodiment, a scope may be realized with the formation of a corresponding peer group. Peer group boundaries may define the search scope when searching for a group's content. For example, a peer in San Francisco looking to buy a used car is normally not interested in cars available outside of the Bay Area. In this case, the peer may want to multicast a message to a subset of the current worldwide peer group, and a subgroup may be formed especially for this purpose. In one embodiment, the multicast may be done without the formation of a new peer group.

10 In one embodiment, all messages may carry a special scope field, which may indicate the scope for which the message is intended. Any peer who receives this message may propagate the message based on the scope indicator. Using this approach, it is preferable that a sending peer is bootstrapped with some well-defined scopes and also has the ability to discover additional scopes.

15 **[0089]** Peer groups may also be formed based upon the proximity of the member peers. Proximity-based peer groups may serve to subdivide the network into abstract regions. Regions may serve as a placeholder for general communication and security configurations that deal with existing networking infrastructure, communication scopes and security requirements. Peer groups may provide a scoping mechanism to reduce traffic overload.

20 **[0090]** Peer groups may provide a secure cooperative environment. Peer group boundaries permit member peers to access and publish protected contents. Peer groups form virtual secure regions which boundaries limit access to the peer group resources. Secure services may be provided to peers within a secured peer group. Their boundaries may or may not reflect any underlying physical network boundaries such as those imposed by routers and firewalls. The concept of a region may virtualize the notion of routers and firewalls, subdividing the network into secure regions in a self-organizing fashion without respect to actual physical network boundaries.

25 **[0091]** Peer groups may also create a monitoring environment. Peer groups may permit peers to monitor a set of peers for any special purpose (heartbeat, traffic introspection, accountability, etc.). Peer groups may also provide a controlled and self-administered environment. Peer groups may provide a self-organized structure that is self-managed and that may be locally managed.

30 **[0092]** Peer groups using the peer-to-peer platform preferably provide several capabilities including, but not limited to, the ability to, find nearby peers, find named peers anywhere on the network, find named peer groups anywhere on the network, join and resign from a peer group, establish pipes between peer group members and find and exchange shared content.

Content

35 **[0093]** Peers may be grouped into peer groups to share content. A content is published and shared among the peer members of a peer group. In one embodiment, content may be shared among group members, but not between groups. In this embodiment, no single item of content may belong to more than one group. If the same content is published in two different peer groups, two different contents may be created. In one embodiment, a content item may be published to make the item's existence known and available to group members through the use of advertisements.

40 **[0094]** An instance of content is a copy of a content. Each content copy may be replicated on different peers in the peer group. Each copy preferably has the same content identifier as well as a similar value. Replicating contents within a peer group may help any single item of content be more available. For example, if an item has two instances residing on two different peers, only one of the peers needs to be alive and respond to the content request. In one embodiment, the peer-to-peer platform protocols do not specify how or when contents are replicated. In one embodiment, whether and how to copy an item of content may be a policy decision that may be encapsulated in higher-level applications and services, for example a content management service.

45 **[0095]** A content may be any computer content (e.g. code, data, applications, active content such as services, or other collection of computer-representable resources). Examples of content include, but are not limited to, a text file, a structured document (e.g. a PDF or a XML file), a Java "jar" or loadable library, code or even an executable process (checkpointed state). No size limitation is assumed. Each content instance may reside on a different peer in the peer group. The instances may differ in their encoding type. HTML, XML and WML are examples of encoding types. Each instance may have the same content identifier as well as a similar set of elements and attributes, and may even exist on the same peer. An encoding metadata element may be used to differentiate instances of content. Making new instances of content on different peers may help any single item of content be more available. For example, if an item

50 has two instances residing on two different peers, only one of the peers needs to be alive and respond to the content request.

55 **[0096]** Items of content that represent a network service may be referred to as active content. These items may have additional core elements above and beyond the basic elements used for identification and advertisement. In one em-

bodiment, active content items may be recognized by Multi-Purpose Internet Mail Extensions (MIME) content type and subtype. In one embodiment, all peer-to-peer platform active contents may have the same type. In one embodiment, the subtype of an active content may be defined by network service providers and may be used to imply the additional core elements belonging to active content documents. In one embodiment, the peer-to-peer platform may give latitude to service providers in this regard, yielding many service implementation possibilities.

[0097] In one embodiment, each item of content may have a unique canonical name. Figure 3 illustrates an exemplary canonical content name (which may be referred to as a content identifier or content ID) according to one embodiment. The unique identifier may include a peer group universal unique identifier (UUID) 170, and also may include another name 174 that may be computed, parsed, and maintained by peer group members. In one embodiment, the UUID may be a 128-bit field. In one embodiment, the name may be a byte array. In one embodiment, the particular name implementation within a peer group is not mandated by the peer-to-peer platform. The name may be, for example, a hash code, a URI, a URN, or a name generated by any suitable means of uniquely identifying content within a peer group. In one embodiment, a length of remainder field 172 may specify the length of the name field 174 for this content in this particular implementation.

[0098] In one embodiment, once a content item has been published to the peer-to-peer network, it may not be assumed that that the content can be later retrieved from the network. The content may be only available from peers that are not currently reachable or not currently part of the network. In one embodiment, once a content item has been published to the peer-to-peer network, it may not be assumed that the content can be deleted. Replication/republishing of content by peers on the network may be unrestricted and the content may propagate to peers that are not reachable from the publishing peer.

Pipes

[0099] Pipes may provide the primary channels for communication among peers and are a mechanism for establishing communication between peers. Pipes may be used as communication channels for sending and receiving messages between services or applications over peer endpoints. Pipes may connect peers that have a direct physical link and peers that do not have a direct physical link. In the latter case, one or more intermediary peer endpoints may be used to route messages between the two pipe endpoints. A pipe instance is, logically speaking, a resource within a peer group. The actual implementation of a pipe instance is typically through a pipe service. In one embodiment, at each endpoint, software to send, or receive, as well as to manage optional associated pipe message queues is assumed, but not mandated.

[0100] Pipes in the peer-to-peer platform are preferably asynchronous, unidirectional, stateless and unreliable to provide the lowest overhead. Pipes are preferably unidirectional, and thus in one embodiment there are input pipes and output pipes. Asynchronous pipes may enable developers to build large-scale interconnected distributed services and applications. Pipes are preferably indiscriminate and may thus support binary code, data strings, Java technology-based objects, and/or applets, among others. The peer-to-peer platform preferably does not define how the internals of a pipe work. Any number of unicast and multicast protocols and algorithms, and combinations thereof, may be used. In one embodiment, one pipe may be chained together with each section of the chain using a different transport protocol.

[0101] The pipe endpoints may be referred to as input pipes (receiving end) and output pipes (sending end). Pipes may provide the illusion of a virtual in and out mailbox that is independent of any single peer location. Services and applications may communicate through pipes without knowing on which physical peer a pipe endpoint is bound. When a message is sent into a pipe, the message is sent to all peer endpoints currently connected (listening) to the pipe. The set of currently connected pipe endpoints (input pipes) may be obtained using the pipe binding protocol.

[0102] Unlike conventional mechanisms, peer-to-peer platform pipes may have ends that may be moved around and bound to different peers at different times, or not connected at all. In one embodiment, pipes may be virtual, in that a pipe's endpoint may be bound to one or more peer endpoints. In one embodiment, pipe endpoints may be non-localized to a physical peer, and may be dynamically bound at creation time or runtime via the pipe binding protocol. The pipe binding process may include discovering and connecting the two or more endpoints of a pipe.

[0103] Using pipes, developers may build highly available services where pipe connections may be established independently of a peer location. This dynamic binding of pipes helps to provide redundant implementation of services over a P2P network. A peer may logically "pick up" a pipe at any point in time. For example, a peer that wants to use a spell checker service may connect to a peer group's spell checker pipe that is implemented as a redundant peer group service. The peer may be serviced as long as there is at least one single instance of a spell checker service still running somewhere within the peer group. Thus, using pipes as described herein, a collection of peers together may provide a high level of fault tolerance, where a new peer at a different location may replace a crashed peer, with the new peer taking over the existing pipe to keep the communication going.

[0104] In one embodiment, enhanced pipes with additional properties such as reliability, security, and quality of service may be supported. In embodiments where the peer-to-peer platform runs on top of transports that have such prop-

erties, an implementation may optimize and utilize the transports. For example, when two peers communicate with each other and both have TCP/IP support, then an implementation may use the bidirectional capabilities of TCP/IP to create bidirectional pipes. Other data transfer methods that may be implemented by pipes as provided at the service layer to provide different quality of service include, but are not limited to: synchronous request-response (the endpoint sends a message, and receives a correlated answer), streaming (efficient control-flow data transfer) and bulk transfer (bulk reliable data transfer of binary data).

[0105] Pipes may offer several modes of communication. Figure 4 illustrates a point-to-point pipe connection between peers 200C and 200D according to one embodiment. In one embodiment, a point-to-point pipe connects exactly two peer endpoints together, an input pipe 202A that receives messages sent from an output pipe 204A. The pipe appears as an output pipe to the sender and as an input pipe to the receiver, with traffic going in one direction only — from the sender to the receiver. In one embodiment, no reply or acknowledgement operation is supported. In one embodiment, additional information in the message payload (for example, a unique ID) may be required to thread message sequences. The message payload may also contain a pipe advertisement that can be used to open a pipe to reply to the sender (send/response).

[0106] Figure 4 also illustrates a propagate pipe with peer 200A as a propagation source and peers 200B and 200C with listening input pipes according to one embodiment. A propagate pipe may connect two or more peer endpoints together, from one output pipe 204B to one or more input pipes (e.g. 202B and 202C). The result is that any message sent into the output pipe is sent to all input pipes. Messages flow into the input pipes from the output pipe (propagation source). A propagate message may be sent to all listening input pipes. This process may create multiple copies of the message to be sent. On transports that provide multicast (e.g. TCP/IP), when the propagate scope maps to underlying physical subnets in a one-to-one fashion, transport multicast may be used as an implementation for propagate. Propagate may be implemented using point-to-point communication on transports that do not provide multicast such as HTTP.

Messages

[0107] In one embodiment, the peer-to-peer platform may use asynchronous messages as a basis for providing Internet-scalable peer-to-peer communication. The information transmitted using pipes may be packaged as messages. Messages define an envelope to transfer any kinds of data. A message may contain an arbitrary number of named subsections which can hold any form of data. In one embodiment, the messages may be in a markup language. In one embodiment, the markup language is XML. Each peer's messaging layer may deliver an ordered sequence of bytes from the peer to another peer. The messaging layer may send information as a sequence of bytes in one atomic message unit. In one embodiment, messages may be sent between peer endpoints. In one embodiment, an endpoint may be defined as a logical destination (e.g. embodied as a URN) on any networking transport capable of sending and receiving Datagram-style messages. Endpoints are typically mapped into physical addresses by the messaging layer at runtime.

[0108] In one embodiment, a message may be a Datagram that may include an envelope and a stack of protocol headers with bodies and an optional trailer. The envelope may include, but is not limited to, a header, a message digest, (optionally) the source endpoint, and the destination endpoint. In one embodiment, each protocol header may include, but is not limited to, a tag naming the protocol in use and a body length. Each protocol body may be a variable length amount of bytes that is protocol tag dependent. Each protocol body may include, but is not limited to, one or more credentials used to identify the sender to the receiver. Such a message format preferably supports multiple transport standards. An optional trailer may include traces and accounting information.

[0109] The messaging layer may use the transport specified by the URN to send and receive messages. In one embodiment, both reliable connection-based transports such as TCP/IP and unreliable connectionless transports like UDP/IP may be supported. Other existing message transports such as IRDA, and emerging transports like Bluetooth may also be supported using the peer endpoint addressing scheme. Peer-to-peer platform messages are preferably useable on top of asynchronous, unreliable, and unidirectional transport. The peer-to-peer platform protocols preferably use a low-level message transport layer (e.g. XML) as a basis for providing Internet-scalable peer-to-peer communication. The peer-to-peer platform preferably does not assume that the networking transport is IP-based.

[0110] The message digest in the envelope may be used to guarantee the data integrity of messages. Messages may also be encrypted and signed for confidentiality and refutability. In one embodiment, each protocol body may include one or more credentials used to identify the sender to the receiver. A credential is a key that, when presented in a message body, may be used to identify a sender and to verify that sender's right to send the message to the specified endpoint. The credential may be an opaque token that is preferably presented each time a message is sent. In one embodiment, the sending address placed in the message envelope may be crosschecked with the sender's identity in the credential. Credentials may be stored in the message body on a per-protocol <tag> basis. In one embodiment, the exact format and content of the credentials are not specified by the peer-to-peer platform. For example,

a credential may be a signature that provides proof of message integrity and/or origin. As another example, a message body may be encrypted, with the credential providing further information on how to decrypt the content. In one embodiment, each credential's implementation may be specified as a plug-in configuration, which preferably allows multiple authentication configurations to coexist on the same network.

[0111] When an unreliable networking transport is used, each message may be delivered more than once to the same destination or may not arrive at the destination. Two or more messages may arrive in a different order than sent. In one embodiment, high-level communication services layered upon the core protocols may perform message re-ordering, duplicate message removal, and processing acknowledgement messages that indicate some previously sent message actually arrived at a peer. Regardless of transport, messages may be unicast (point to point) between two peers or may be propagated (like a multicast) to a peer group. Preferably, no multicast support in the underlying transport is required. In one embodiment, peers receiving a corrupted or compromised message may discard the message. Messages may be corrupted or intentionally altered in transmission on the network.

[0112] The peer-to-peer platform preferably does not mandate how messages are propagated. For example, when a peer sends out a peer discovery message, the peer discovery protocol preferably does not dictate if the message should be confined to the local area network only, or if it must be propagated to every corner of the world.

[0113] The peer-to-peer platform messages 252 are preferably defined with the envelope 250 as illustrated in Figure 5. In one embodiment, the messages are defined in a markup language. In one embodiment, the markup language is XML. The following is an exemplary message in XML:

```
<?xml version="1.0" encoding="UTF-8"?>
  <SampleMessage>
    <SampleMessageVersion> version number "1.0" </SampleMessageVersion>
    <SampleMessageDest> destination peer id </SampleMessageDest>
    <SampleMessageSrc> source peer id </SampleMessageSrc>
    <SampleMessageDigest> digest </SampleMessageDigest>
    <SampleMessageTagName> tag </SampleMessageTagName>
    <SampleMessageTagData> body </SampleMessageTagData>
    .....
    <SampleMessageTagName> tag </SampleMessageTagName>
    <SampleMessageTagData> body </SampleMessageTagData>
    .....
    <SampleMessageTrailer> String</ SampleMessageTrailer >
  </SampleMessage>
```

[0114] The version number may be a string. The destination and source peer identifier may be represented as peer-to-peer platform identifiers. In one embodiment, the digest is either an MD5 or SHA1 hash or a digital signature. The digest may serve as a placeholder for either. A message may have as many tag parts as needed. In one embodiment, the tag name may be a string and the body may be a byte array containing a string without XML escape characters ("<", ">") or a base64 encoded string.

[0115] In one embodiment, the message format may support binary data and/or multi-part messages with MIME-types. The message format may allow for arbitrary message header fields, including optional header fields. The message format may allow for data verification of message content and the cryptographic signing of messages. The message format may provide an arbitrary number of named subsections that may contain any form of data of any (reasonable) size. The message format may be "email-safe" such that its contents may be extracted reliably after standard textual transformations committed my E-mail client and server software.

Services

[0116] Peers may cooperate and communicate to publish, discover and invoke network services. A service denotes a set of functions that a provider offers. In one embodiment, a peer-to-peer platform peer can offer a service by itself or in cooperation with other peers. In one embodiment, a peer may publicize a service by publishing a service adver-

tisement for the service. Other peers may then discover this service using the peer discovery protocol (through the advertisement) and make use of it. A peer may publish as many services as it can provide.

[0117] In one embodiment, services may either be pre-installed into a peer or loaded from the network. The process of finding, downloading and installing a service from the network may include performing a search on the network for the service, retrieving the service, and then installing the service. Once a service is installed and activated, pipes may be used to communicate with the service. In one embodiment, peer-to-peer platform-enabled services may publish pipe advertisements as their main invocation mechanism. The service advertisement may specify one or more pipe advertisements that may be used by a peer to create output pipes to invoke the service. The service advertisement may also include a list of predetermined messages that may be sent by a peer to interact with the service. The service advertisement may describe all messages that a client may send or receive.

[0118] Several methods may be provided by various embodiments to publish a service. Services may be published before creating a new peer group by adding the service advertisement to the peer group advertisement. Services may also be published by adding the services in a separate peer service advertisement. The discovery service may also allow new advertisements to be added at runtime. The new advertisement will belong to a predefined peer group. Other methods of publishing services may be provided. Note that service advertisements may be placed in the peer group advertisement of any group. Since all peers belong to the global peer group, a peer may publish the service in the global peer group advertisement to make it available to any peer.

[0119] In one embodiment, services advertised in a peer group advertisement are instantiated for a peer when the peer joins the group. In one embodiment, all the services are instantiated. In another embodiment, none, one or more of the advertised services may be instantiated when the peer joins the peer group. Service advertisements in the peer group advertisement may include resolver, discovery, membership, peer information and pipe service advertisements. In one embodiment, services advertised in a peer group advertisement are loaded on the peer when the peer boots. In one embodiment, this automated loading is not mandatory but is part of the Java Binding. One embodiment may provide a mechanism to force a service in a peer group advertisement to be instantiated by a peer.

[0120] In one embodiment, when a peer boots, any services advertised in the peer advertisement are loaded. The peer advertisement corresponds to the platform advertisement. These services may include the minimal set of services to bootstrap the creation of new peers: discovery service, membership service, resolver service, peer information service and pipe service.

[0121] In one embodiment, when a peer switches from one peer group to another, the first group's services remain active. In one embodiment, a peer may call a stop method on the service application interface to stop an instance of a local service. A peer that is a member of one peer group that refers to a service may join a second peer group that also refers to the service while still a member of the first. Whether the service is instantiated once or twice may depend on the service implementation. Some service implementations may use a static instantiation that is done once. In this case, all groups share the same instance. Other service implementations are local to a peer group and are not aware of the state of any other peer groups on the same node.

[0122] In one embodiment, services may use a "time to live" indicator that defines when the service was created, and would also define the lifetime of the service. After its lifetime has expired, the stale service may be purged.

[0123] A service may be well defined and widely available so that a peer can use it directly. Other services may require special code in order to actually access the service. For example, the way to interface with the service provider may be encoded in a piece of software. In this case, it is preferable if a peer can locate an implementation that is suitable for the peer's specific runtime environment. In one embodiment, if multiple implementations of the same service are available, then peers hosted on Java runtimes can use Java programming language implementations while native peers to use native code implementations. In one embodiment, service implementations may be pre-installed into a peer node or loaded from the network. In one embodiment, once a service is installed and activated, pipes may be used to communicate with the service.

[0124] In one embodiment, each service may have a unique identifier. In one embodiment, a service may have a name that may include a canonical name string that may indicate the type and/or purpose of the service. A service may also provide optional information (e.g. a set of descriptive keywords) that further describes the service. The unique identifier, name and optional information may be stored within a service advertisement. The advertisement may also include other information needed to configure and instantiate a service.

[0125] In one embodiment, the peer-to-peer platform may recognize two levels of services, peer services and peer group services. A service that executes only on a single peer may be referred to as a peer service. A peer service is accessible only on the peer that is publishing the service. If that peer happens to fail, then service also fails. This level of service reliability may be acceptable for an embedded device, for example, providing a calendar and email client to a single user. Multiple instances of the service may be run on different peers, but each instance publishes its own advertisement. A service that is composed of a collection of cooperating instances (potentially cooperating with each other) of the service running on multiple peers in a peer group may be referred to as a peer group service. A peer group service may employ fault tolerance algorithms to provide the service at a higher level of availability than that a

peer service can offer. If any one peer fails, the collective peer group service may not be affected, because the service may still be available from at least one other peer member. Peer group services may be published as part of the peer group advertisement.

[0126] In one embodiment, the peer-to-peer platform may include a set of default peer group services such as peer discovery, as well as a set of configurable services such as routing. In one embodiment, a peer-to-peer platform peer may not be required to have one or all of these services. For example, a cell phone peer may be pre-configured with enough information to contact a fixed server provided by the telecom operator. This may be enough to bootstrap the cell phone peer without requiring it to independently carry with it additional services.

[0127] In one embodiment, although the concept of a service is orthogonal to that of a peer and a peer group, a peer group formed using the peer-to-peer platform may require a minimum set of services needed to support the operation of the group. Some services may be well known and may be referred to as peer-to-peer platform core services. Embodiments of the peer-to-peer platform may define a set of core peer group services that may be used to form and support peer groups. In one embodiment, the core peer group services may provide the minimum services required to form a peer group (e.g. membership and discovery services). The peer-to-peer platform core services are preferably 100% decentralized and thus may enable pure peer-to-peer network computing. In one embodiment, it is not required that all core services be implemented by every peer group.

[0128] In one embodiment, the peer-to-peer platform may define peer group core services including, but not limited to, a discovery service, a membership service, an access service, a pipe service, a resolver service and a monitoring service. A discovery service may be used to search for peer group resources such as peers, peer groups, and pipes. The search criteria may include a resource name. Discovery and discovery services are described more fully later in this document.

[0129] In one embodiment, most peer groups will have at least a membership service. Current peer group members may use the membership service during the login process to reject or accept a new peer group membership application.. The membership service may be a "null" authenticator service that imposes no real membership policy. Peers wishing to join a peer group first locate a current member, and then request to join. The application to join may be either rejected or accepted by the collective set of current members. The membership service may enforce a vote of peers or alternatively elect a designated group representative to accept or reject new membership applications.

[0130] An access service may be used to validate, distribute, and authenticate a group member's credentials. The access service may define the type of credential used in the message-based protocols used within the peer group. The access service may be used to validate requests made by one peer to another. The peer receiving the request provides the requesting peer's credentials and information about the request being made to the access service to determine if the access is permitted. In one embodiment, not all actions within the peer group need to be checked with the access service, only those actions which only some peers are permitted to use.

[0131] A pipe service may be used to manage and create pipe connections between the different peer group members. A resolver service may be used to send query string to peers to find information about a peer, a peer group, a service or a pipe. A monitoring service is used to allow one peer to monitor other members of the same peer group.

[0132] In one embodiment, not all the above services are required to be implemented by a peer group. Each service may implement one or more of the peer-to-peer platform protocols. A service preferably implements one protocol for simplicity and modularity reasons, but some services may not implement any protocols.

[0133] Other services may be user-defined and provide application dependent services such as content searching and indexing. A user-defined service may provide additional APIs. User-defined services may be implemented that may offer the ability to mix-in centralization as a means of increasing performance. In one embodiment, the peer-to-peer platform core services may provide a reference implementation for user-defined services. Examples of user defined services may include, but are not limited to:

- Efficient long-distance peer lookup and rendezvous using a peer naming and discovery service.
- Simple, low-cost information search and indexing using a content sharing service.
- Interoperability with existing centralized networking infrastructure and security authorities in corporate, public, private, or university networks using administration services.
- A resolver service may be implemented to find active (running on some peer) and inactive (not yet running) service instances.
- An FTP service that allows file transfers among peers over pipes using FTP.

Advertisements

[0134] In one embodiment, the peer-to-peer protocols may use advertisements to describe and publish the existence of peer resources. An advertisement may be defined as a structured, language neutral metadata structure that names, describes, and publishes the existence of a peer-to-peer platform resource, such as a peer, a peer group, a pipe, or

a service.

[0135] In one embodiment, advertisements may be used in the peer-to-peer platform as language-neutral metadata structures. In one embodiment, each software platform binding may describe how advertisements are converted to and from native data structures such as Java objects or 'C' structures. Each protocol specification may describe one or more request and response message pairs. In one embodiment, advertisements may be the most common document exchanged in messages.

[0136] Information exchanged between peers may include advertisement documents. The peer-to-peer platform preferably includes advertisement documents to represent all of the peer-to-peer platform resources managed by the core platform, such as peers, peer groups, pipes and services. In one embodiment, the peer-to-peer platform may define a set of core advertisements. The peer-to-peer platform may define core advertisement types including, but not limited to, one or more of peer advertisements, peer group advertisements, pipe advertisements, service advertisements, content advertisements, and endpoint advertisements. In one embodiment, user-defined advertisement subtypes (for example, using XML schemas) may be formed from these basic types. Subtypes of the core advertisements may be used to add an unlimited amount of extra, richer metadata to a peer-to-peer network. The peer-to-peer platform protocols, configurations and core software services however, preferably operate only on the core advertisements.

[0137] In one embodiment, an advertisement is a markup language structured document that names, describes, and publishes the existence of a peer-to-peer platform resource. In one embodiment, peer-to-peer platform advertisements may be represented in the Extensible Markup Language (XML) and are therefore software platform neutral. XML provides a powerful means of representing data and metadata throughout a distributed system. XML provides universal (software-platform neutral) data because XML is language agnostic, self-describing, strongly-typed and ensures correct syntax. XML advertisements may be strongly typed and validated using XML schemas. XML also allows advertisements to be translated into other encodings such as HTML and WML. This feature allows peers that do not support XML to access advertised resources. In one embodiment, each document may be converted to and from a platform specific representation such as a Java object. In one embodiment, peers supporting the various protocols requiring that advertisements be exchanged in messages may accept only valid XML documents that descend from the base XML advertisement types.

[0138] Advertisements represented in a markup language such as XML, like any markup language document, may be composed of a series of hierarchically arranged elements. Each element may include its data and/or additional elements. An element may also have attributes. Attributes are name-value string pairs. An attribute may be used to store metadata, which may be used to describe the data within the element.

[0139] In one embodiment, a peer advertisement may be used to describe a peer. A peer advertisement may describe the peer resources. One use of a peer advertisement is to hold specific information about the peer, such as its name, peer identifier, registered services and available endpoints. Figure 6 illustrates the content of a peer advertisement according to one embodiment. The following is an example of one embodiment of a peer advertisement in XML, and is not intended to be limiting:


```

5      <?xml version="1.0" encoding="UTF-8"?>
      <PeerAdvertisement>
          <Name> name of the peer</Name>
          <Keywords>search keywords </Keywords>
          <Pid> Peer identifier </Pid>
10      <Services>
          < Service advertisement>
              ....
          </Service advertisement>
15      </Services>
          <Endpoints>
          <endpoint advertisement >
20      ....
          </endpoint advertisement >
          </Endpoint>
          <InitialApp>
25      < Service advertisement >
          ....
          </ Service advertisement >
30      </InitialApp>
      </PeerAdvertisement>

```

[0140] Embodiments of a peer advertisement may include, but are not limited to, the following fields:

- 35 • Name: an optional string that can be associated with a peer. In one embodiment, the name is not required to be unique unless the name is obtained from a centralized naming service that guarantees name uniqueness.
- Keywords: an optional string that may be used to index and search for a peer. In one embodiment, the string is not guarantee to be unique. Two peers may have the same keywords. The keywords string may contain spaces.
- 40 • Peer identifier: uniquely identifies the peer. In one embodiment, this may be a required element. Each peer has a unique identifier.
- Service: a service advertisement element for each service published on the peer. Services started on a peer may publish themselves to the peer. In one embodiment, not all services running on the peer need to publish themselves.
- 45 • Endpoint: an endpoint URI (e.g. tcp://129.144.36.190:9701 or http://129.144.36.190:9702) for each endpoint available on the peer.
- InitialApp: Optional application/service started when the peer is booted. A service advertisement is used to describe the service.

50 **[0141]** In one embodiment, a peer group advertisement may be used to describe, for a peer group, the group specific information (name, peer group identifier, etc.), the membership process, and the available peer group services. The peer group advertisement defines the core set of services to be used by that peer group. In one embodiment, it may not enforce that each peer must run each service locally. Rather it defines the set of services that are made available to the peer group.

55 **[0142]** In one embodiment, the initial creator of the peer group may define what advertisements go into the peer group advertisement at creation time. Other peers may get a copy of the peer group advertisement when they discover advertisements via the discovery service. In one embodiment, peer group advertisements are immutable objects and new services may not be added due to java binding. Other embodiments may allow new services to be added. In one embodiment, a peer group may provide a registration service that allows the dynamic registration of services.

[0143] Figure 7 illustrates the content of a peer group advertisement according to one embodiment. The following is an example of one embodiment of a peer group advertisement in XML, and is not intended to be limiting:

```

5      <?xml version="1.0" encoding="UTF-8"?>
      <PeerGroupAdvertisement>
        <Name> name of the peer group</Name>
        <Keywords>search keywords</Keywords>
10     <Gid> Peer group Id</Gid>
        <Services>
          <Service advertisement>
            ...
15         </Service advertisement>
        </Services>
        <InitialApp>
          <Service advertisement>
20         ...
          </Service advertisement>
        </InitialApp>
25     </PeerGroupAdvertisement>

```

[0144] Embodiments of a peer group advertisement may include, but are not limited to, the following fields:

- Name: an optional name that may be associated with a peer group. In one embodiment, the name is not required to be unique unless the name is obtained from a centralized naming service that guarantee name uniqueness.
- Keywords: an optional string that may be used to index and search for a peer group. In one embodiment, the string is not guarantee to be unique. Two peer groups may have the same keywords.
- Peer group Id: uniquely identifies the peer group. In one embodiment, this is a required element. Each peer group has a unique id.
- Service: a service advertisement element for each peer group service available in the peer group. In one embodiment, not all peer group services need to be instantiated when a peer joins a peer group. In one embodiment, at least a membership service should be available, so the membership service may implement a null authenticator membership.
- InitialApp: optional application/service started when a peer is joining a peer group. A service advertisement may be used to describe the service. The initial application may be started when a peer is joining a group. Alternatively, it may be left to the joining peer to decide to either start or not start the peer group initial application.

[0145] Once a peer joins a group, that peer may receive (depending again upon membership configuration) a full membership-level peer group advertisement. The full membership advertisement, for example, might include the configuration (required of all members) to vote for new member approval.

[0146] In one embodiment, a pipe advertisement may be used to describe an instance of a pipe communication channel. A pipe advertisement may be used by a pipe service to create associated input and output pipe endpoints. In one embodiment, a pipe advertisement document may be published and obtained either by using a discovery service (e.g. the core discovery service) or by embedding it within other advertisements such as the peer or peer group advertisement. Each pipe advertisement may include an optional symbolic name that names the pipe and a pipe type to indicate the type of the pipe (point-to-point, propagate, secure, etc). Figure 8 illustrates the content of a pipe advertisement according to one embodiment. The following is an example of one embodiment of a pipe advertisement in XML, and is not intended to be limiting:

```

<?xml version="1.0" encoding="UTF-8"?>
<PipeAdvertisement>
  <Name> name of the pipe</Name>
  <Id> Pipe Id </Id>
  <Type> Pipe Type </Type>
</PipeAdvertisement>

```

[0147] Embodiments of a pipe advertisement may include, but are not limited to, the following fields:

- Name: an optional name that may be associated with a pipe. In one embodiment, the name is not required to be unique unless the name is obtained from a centralized naming service that guarantee name uniqueness.
- Pipe identifier: uniquely identifies the pipe. In one embodiment, this is a required element. Each pipe has a unique id.
- Type: This is an optional pipe type that may be provided to specify the quality of services implemented by the pipe. Pipe types may include, but are not limited to:
 - RELIABLE (guaranteed delivery and ordering, and deliver only once)
 - UNRELIABLE (may not arrive at the destination, may be delivered more than once to the same destination, may arrive in different order)
 - SECURE (reliable and encrypted transfer)

[0148] In one embodiment, a service advertisement may be used to describe a peer-to-peer platform-enabled service. Service advertisements preferably describe how to activate and/or use the service. In one embodiment, a peer-to-peer platform-enabled service is a service that uses pipes as primary invocation mechanism. To invoke the service, a peer may send a message to the associated service pipe. In one embodiment, the core peer group services that each peer group preferably implements in order to respond to the messages described for the peer-to-peer platform protocols are peer-to-peer platform-enabled services and thus may be published using service advertisements. The service advertisement document may be published and obtained using the peer information protocol for peer services, or alternatively using the peer group discovery protocol for peer group services.

[0149] In one embodiment, a pipe advertisement and access method fields may provide a placeholder for any kind of service invocation schema that defines the valid set of XML messages accepted by the service and the associated message flow. Thus, the peer-to-peer platform protocols may be agnostic of service invocation and interoperate with any existing framework. A service advertisement access method field may refer to a WSDL (e.g. www.w3.org/TR/wsdl/), ebXML (e.g. www.ebxml.org/), UPnP (e.g. www.upnp.org/) or a client-proxy schema. among others. For example, a WSDL access method may define messages that are abstract descriptions of the data being exchanged and the collections of operations supported by the service using a WSDL schema. In one embodiment, a service advertisement may include multiple access method tags, as there may be multiple ways to invoke a service. Thus, the peer may ultimately decide which invocation mechanism to use. For example, small devices may want to use a small-footprint mechanism or a service framework they already have the code for, and larger devices may decide to download a client-proxy code.

[0150] In one embodiment, the access method for services is a schema of valid XML messages accepted by the service. In one embodiment, a service advertisement may contain a URL or URI tag to point to a jar file, DLL, or loadable library. A peer may use this to download the code to run the service, for example if the peer joins the peer group and doesn't have the required code to run the service.

[0151] In one embodiment, once a service advertisement is sent out into the world there is no method of pulling it back in. However, each individual peer may have the ability to purge the set of cached advertisements that reside locally, and a rendezvous peer may purge its cache periodically (e.g. daily).

[0152] Figure 9 illustrates the content of a service advertisement according to one embodiment. The following is an example of one embodiment of a service advertisement in XML, and is not intended to be limiting:

```

5      <?xml version="1.0" encoding="UTF-8"?>
      <ServiceAdvertisement>
        <Name> name of the Service</Name>
        <Version> Version Id </Version>
        <Keywords>search keywords </Keywords>
        <Id> Service identifier </Id>
        <Pipe> Pipe endpoint to access the service </Pipe>
10     <Params> service configuration parameters </Params>
        <URI> service provider location</URI>
        <Provider> Service Provider</Provider>
15     <AccessMethods>
        ...
        </AccessMethods>
20     </ServiceAdvertisement>

```

[0153] Embodiments of a service advertisement may include, but are not limited to, the following fields:

- 25 • Name: an optional name that may be associated with a service. In one embodiment, the name is not required to be unique unless the name is obtained from a centralized naming service that guarantees name uniqueness.
- Keywords: an optional string that may be used to index and search for a service. In one embodiment, the string is not guaranteed to be unique. Two services may have the same keywords.
- 30 • Service Id: uniquely identifies a service. In one embodiment, each service has a unique id. In one embodiment, this element may be required.
- Version: specifies the service version number. In one embodiment, this element may be required.
- Provider: gives information about the provider of the service. This will typically be a vendor name. In one embodiment, this element may be required.
- 35 • Pipe: an optional element that specifies a pipe advertisement to be used to create an output pipe to connect to the service. In one embodiment, services are not required to use pipes.
- Params: a list of configuration parameters available to the peer when invoking the service. In one embodiment, the parameter field is optional. Parameters may be defined as a list of strings.
- URI: This is an optional parameter that may be used to specify the location of where the code for the service may be found.
- 40 • Access Methods: In one embodiment, at least one access method is required to specify how to invoke the service. Multiple access method tags may be used when multiple access methods are available. The access method tag allows any kind of service invocation representation to be specified. For example the access method may be a placeholder for a WSDL or uPnP document that describes a web service access method.

45 **[0154]** In one embodiment, a content advertisement may be used to describe a content document stored somewhere in a peer group. In one embodiment, there are no restrictions on the type of contents that can be represented. A content may be a file, a byte array, code or process state, for example. In one embodiment, each item of content may have a unique identifier also known as its canonical name. The unique identifier may include a peer group universal unique identifier (UUID), and also may include another name that may be computed, parsed, and maintained by peer group members. In one embodiment, the content's name implementation within the peer group is not mandated by the peer-to-peer platform. The name may be a hash code, a URI, or a name generated by any suitable means of uniquely identifying content within a peer group. The entire canonical content name may be referred to as a content identifier or content ID. Figure 3 illustrates an exemplary content identifier according to one embodiment.

50 **[0155]** Figure 10 illustrates a content advertisement according to one embodiment. A size element is preferably provided for all content items and gives the total size of the content. In one embodiment, the size is in bytes. In one embodiment, the size is a long (unsigned 64-bits). A content advertisement may also include a MIME (Multi-Purpose Internet Mail Extensions) type that describes the MIME type (encoding may be deduced from the type) of the in-line or referenced data. A content advertisement may also include a RefID element. If the advertised content is another

advertisement (based upon its type), the RefID is the content ID of the referenced content. If the advertised content is not another advertisement, the RefID element may be omitted.

[0156] The following is an example of one embodiment of a content advertisement in XML, and is not intended to be limiting:

```

5      <?xml version="1.0" encoding="UTF-8"?>
      <ContentAdvertisement>
10         <Mimetype> name of the pipe</Mimetype>
          <Size> Pipe Id </Size>
          <Encoding> Pipe Type </Encoding>
          <ID> Content ID</ID>
15         <RefID> Content ID </RefID>
          <Document> document </Document>
      </ContentAdvertisement>
20

```

[0157] Embodiments of a content advertisement may include, but are not limited to, the following fields:

- ID: in one embodiment, all contents have a unique id.
- Size: the total size of the content. In one embodiment, a long (unsigned 64-bits) represented as a string. "-1" indicates that the size is unknown.
- Mimetype: the mime type of the content. The type may be unknown.
- Encoding: specifies the encoding used.
- RefID: if the advertised content is about another content, the RefID specifies the content ID of the referenced content.

[0158] In one embodiment, an endpoint advertisement may be used to describe peer transport protocols. In one embodiment, a peer may support one or more transport protocols. In one embodiment, peers may have multiple network interfaces. Typically, there will be one peer endpoint for each configured network interface and/or protocol (e.g. TCP/IP, HTTP). An endpoint advertisement may be included as a tag field in a peer advertisement to describe the endpoints available on the member peer. In one embodiment, an endpoint advertisement document may be published and obtained either by using the core discovery service or by embedding it within other advertisements such as the peer advertisement. Each endpoint advertisement may include transport binding information about each network interface or transport protocol. Endpoints may be represented with a virtual endpoint address that may include all necessary information to create a physical communication channel on the specific endpoint transport. For example, "tcp://123.124.20.20:1002" or "http://134.125.23.10:6002" are strings representing endpoint addresses. Figure 11 illustrates the content of an endpoint advertisement according to one embodiment. The following is an example of one embodiment of an endpoint advertisement in XML, and is not intended to be limiting:

```

45      <?xml version="1.0" encoding="UTF-8"?>
      <EndpointAdvertisement>
          <Name> name of the endpoint</Name>
          <Keywords> search string </Keywords>
50         <Address> endpoint logical address </Address>
      </EndpointAdvertisement>

```

[0159] Embodiments of an endpoint advertisement may include, but are not limited to, the following fields:

- Name: an optional name that may be associated with an endpoint. In one embodiment, the name is not required to be unique unless the name is obtained from a centralized naming service that guarantee name uniqueness.
- Keywords: an optional string that may be used to index and search for an endpoint. In one embodiment, the string

is not guarantee to be unique. Two endpoints may have the same keywords.

Peer-to-peer Platform Protocols

[0160] The peer-to-peer platform protocols may be used to provide and support ad hoc, pervasive, and multi-hop peer-to-peer (P2P) network computing. Using the protocols, peers can cooperate to form self-organized and self-configured peer groups independently of their positions in the network (e.g. edges, firewalls), and without the need of a centralized management infrastructure. The peer-to-peer platform protocols may have very low overhead, make few assumptions about the underlying network transport and limited requirements of the peer environment, and may be used to deploy a wide variety of P2P applications and services in a highly unreliable and changing network environment.

[0161] In one embodiment, the peer-to-peer platform may include core protocols including, but not limited to, a peer membership protocol, a peer discovery protocol, a peer resolver protocol, a peer information protocol, a pipe binding protocol, and a peer endpoint protocol. These protocols may be implemented using a common messaging layer. This messaging layer binds the protocols to various network transports. In one embodiment, the peer-to-peer platform protocols may be specified as a set of markup language (e.g. XML) messages exchanged between peers. Each software platform binding describes how a message is converted to and from a native data structures such as a Java object or 'C' structure. In one embodiment, the use of markup language messages to define protocols allows many different kinds of peers to participate in a protocol. Each peer is free to implement the protocol in a manner best suited to its abilities and role. Peer-to-peer platform messages are described previously in this document.

[0162] In one embodiment, each of the protocols is independent of the others. Preferably, a peer is not required to implement all of the networking protocols. A peer preferably needs to implement only the protocol that it requires. For example, a device may have all the advertisements it uses pre-stored in memory, so that peer does not need to implement the Peer Discovery Protocol. As another example, a peer may use a pre-configured set of peer routers to route all its messages, hence the peer does not need to implement the Peer Endpoint protocol. Instead, the peer sends messages to the routers to be forwarded. As yet another example, a peer may not need to obtain or wish to provide status information to other peers, hence the peer does not to implement the peer information protocol. The same can be said about all of the other protocols. In one embodiment, a peer may implement only a portion (client-side or server-side only, for example) of a protocol.

[0163] Peers may use the peer-to-peer platform protocols to advertise their resources and to discover network resources (services, pipes, etc.) available from other peers. Peers may form and join peer groups to create special relationships. The peer-to-peer platform protocols may allow peers to communicate without needing to understand or manage the potentially complex and dynamic network topologies that are becoming common. Peers may cooperate to route messages allowing for full peer connectivity. The peer-to-peer platform protocols allow peers to dynamically route messages across multiple network hops to any destination in the network (potentially traversing firewalls). Each message may include either a complete or a partial ordered list of gateway peers through which the message might be routed. If route information is incorrect, an intermediate peer may assist in dynamically finding a new route. A peer-to-peer platform protocol message that is routed through multiple hops is preferably not assumed to be reliably delivered, even if only reliable transports such as TCP/IP are used through all hops. A congested peer may drop messages at any time rather than routing them.

[0164] The peer-to-peer platform protocols may be implemented on a variety of networks including, but not limited to, the Internet, corporate intranets, dynamic proximity networks, home networking environments, LANs and WANs. The peer-to-peer platform protocols may allow the peer-to-peer platform to be easily implemented on unidirectional links and asymmetric transports. In particular, many forms of wireless networking do not provide equal capability for devices to send and receive. The peer-to-peer platform permits any unidirectional link to be used when necessary, improving overall performance and network connectivity in the system. Thus, the peer-to-peer platform protocols may be easy to implement on any transport. Implementations on reliable and bidirectional transports such as TCP/IP or HTTP may provide efficient bidirectional communications. Even on bidirectional transports, communication ability between any pair of peers may at times not work equally well in both directions. That is, communications between two peers will in many cases be able to operate bidirectionally, but at times the connection between two peers may be only unidirectional, allowing one peer to successfully send messages to the other while no communication is possible in the reverse direction. The peer-to-peer platform unidirectional and asymmetric transport also plays well in multi-hop network environments where the message latency may be difficult to predict. Furthermore, peers in a P2P network tend to have nondeterministic behaviors and thus may appear or leave the network very frequently.

[0165] In one embodiment, the peer-to-peer platform protocols do not require a broadcast or multicast capability of the underlying network transport. Messages intended for receipt by multiple peers (propagation) may be implemented using point-to-point communications. The peer-to-peer platform protocols preferably do not require periodic messages of any kind at any level to be sent within the network, and thus preferably do not require periodic polling, link status sensing, or neighbor detection messages, and may not rely on these functions from any underlying network transport

in the network. This entirely on-demand behavior of the protocols and lack of periodic activity may allow the number of overhead messages generated by the peer-to-peer platform to scale all the way down to near or at zero, when all peers are stationary with respect to each other and all routes needed for current communication have already been discovered.

[0166] In one embodiment, the peer-to-peer platform protocols are defined as idempotent protocol exchanges. The same messages may be sent/received more than once during the course of a protocol exchange. Preferably, no protocol states are required to be maintained at both ends. Due to the unpredictability of P2P networks, assumptions may not be made about the time required for a message to reach a destination peer, and thus the peer-to-peer platform protocols preferably do not impose any timing requirements for message receipt.

[0167] The peer-to-peer platform protocols may take advantage of additional optimizations, such as the easy ability to reverse a source route to obtain a route back to the origin of the original route.

[0168] Figure 12 illustrates protocols and bindings in a peer-to-peer platform according to one embodiment. When the peer-to-peer platform protocols are implemented using a particular programming language and over a particular transport protocol, the implementation is an instance of a peer-to-peer platform binding 220, where the peer-to-peer platform protocols are bound to the language and the transport layer. In one embodiment, protocol and peer software implementation issues may be defined in documents specific to the binding. A binding document describes how the protocols are bound to an underlying network transport (such as TCP/IP or UDP/IP) or to a software platform such as Java 222 or a native software platform 224 such as UNIX.

[0169] The following describes the transport binding of the peer-to-peer platform protocols over TCP/IP including the message wire format of peer-to-peer platform endpoint messages over a TCP/IP socket connection according to one embodiment. Each TCP/IP message may include a header and a body. In one embodiment, the format of the header is:

Type	Source IP address	Source Port	Size	Option	Unused
------	-------------------	-------------	------	--------	--------

[0170] The type may include information used to either unicast or multicast the request. The type may indicate whether this is a propagate message, a unicast message, an ACK or a NACK. The port may allow each peer to decide to bind its transport service to a specific port number. The TCP binding preferably does not require that a specific port be used. The size may indicate the body size (not including the header). The option may be used to specify the kind of socket connections (uni- or bi-directional) in use. The TCP/IP binding does not require the maintenance of any states. The normal operation is for one peer to send a TCP/IP packet to another one, and to close the socket after the packet is sent. This is the minimum functionality required to implement unidirectional pipes. In one embodiment, if the receiving end decides to keep the connection active (socket "keep alive"), it may return an indicator to the sender to tell the sending end that it is keeping the connection alive. The sending end may reuse the same socket to send a new packet.

[0171] The following describes the transport binding of the peer-to-peer platform protocols over HTTP including the wire message format for the HTTP binding of the peer-to-peer platform protocols. An HTTP request format message may include a header and a body using an HTML format. For example:

```
<HTML>
  <Code> Header </Code>
  <Msg> Body </Msg>
</HTML>
```

[0172] The header allows the receiving end to determine which message type is received. Message types may include request succeeded, request failed, empty (no body) and response (the body is not empty and contains data). The body may be represented as a string in the HTML request document. Connection states that may be used include, but are not limited to:

- Peer Connection:** Before a message can be sent to a HTTP server peer, the HTTP client may be required to send a request for connection to the other peer. The request for connection message may use the empty header type. The message may be sent using a GET request to the following server URL: `http://ip-name:port/reg/client-peerid/`. `ip-name` specifies the IP of the server peer and the port is the corresponding server port number (8080 for example). The server replies with an empty message containing either a request succeeded or request failed header type. The peer connection message may be used to create a client session on the receiving peer. The receiving peer may decide to reject the connection and refuse the client connection. This corresponds to a client registration.
- Message Sending:** To send a message to another peer server, the client sends a message of the response type with a message body part. The server replies with an ok or failed message. The message is sent to the following URL using the PUT method: `http://ip-name:port/snd/`. The server replies with a message including a request suc-

ceeded or request failed header type.

- **Message Retrieving:** To retrieve messages from a peer server, the client may send a GET request message with the empty header tag to the following URL: `http://ipname:port/rec/client-peerid/`. The server replies with may respond with a message failed message or with a Content message including the messages retrieved.

Peer Discovery Protocol

[0173] In one embodiment, the peer-to-peer platform may include a peer discovery protocol that may allow a peer to find advertisements on other peers. The peer discovery protocol may be used to discover any published peer resources including other peers, peer groups, pipes, services and any other resource that has an advertisement in the peer-to-peer network. This protocol may be used to find members of any kind of peer group, presumably to request membership. In one embodiment, the peer discovery protocol is the default discovery protocol for all peer groups, including the world peer group. The discovery protocol may be used as a default discovery protocol that allows all peer-to-peer platform peers to understand each other at a very basic level.

[0174] The peer discovery protocol may provide, at the lowest level, the minimum building blocks for propagating discovery requests between peers. Thus, the peer discovery protocol may provide the essential discovery infrastructure for building high-level discovery services. In many situations, discovery information is better known by a high-level service, because the service may have a better knowledge of the topology (firewall traversal), and the connectivity between peers. The peer discovery protocol may provide a basic mechanism to discover advertisements while providing hooks so high-level services and applications can participate in the discovery process. Services may be able to give hints to improve discovery (i.e. decide which advertisements are the most valuable to cache).

[0175] In one embodiment, the peer discovery protocol may be based on web crawling and the use of rendezvous peers. Rendezvous peers are peers that offer to cache advertisements to help others peers discover resources, and propagate requests they cannot answer to other known rendezvous peers. Rendezvous peers and their use in the discovery process are discussed later in this document.

[0176] In one embodiment, custom discovery services may choose to leverage the peer discovery protocol. If a peer group does not have its own discovery service, the peer discovery protocol is preferably used as the method for probing peers for advertisements. Rendezvous peers may keep a list of known peers and peer groups. This list may or may not be exhaustive or timely. A custom discovery service (if it knew that the region's rendezvous did keep a timely exhaustive list), for example, may discover all peers in the region by sending a single message to the rendezvous peer.

[0177] In one embodiment, peer discovery may be done with, or alternatively without, specifying a name for the peer to be located and/or the group to which peers belong. When no name is specified, all discovered advertisements of the requested type may be returned. If a probing peer provides the name of the peer to be located, a simple translation may be requested that returns that peer's advertisement. Once a peer is discovered, ping, status, and capability messages may be sent to its "main" endpoint(s) using a peer information protocol. Peers may export more than one endpoint. Preferably, each peer designates at least one primary endpoint to handle the low-level housekeeping protocols such as the peer discovery protocol and the peer information protocol.

[0178] In one embodiment, the peer discovery protocol may be used to probe network peer groups looking for peers that belong to specified peer groups. This process may be referred to as screening. Peers may be screened for membership by presenting each candidate member with a peer group name (string matched with the peer group advertisement canonical name). Preferably, peers claiming to belong to this group may respond, while other peers do not respond. The peer discovery protocol may be used to discover any type of core advertisement including, but not limited to: peer advertisements, peer group advertisements, pipe advertisements and service advertisements.

[0179] Peer groups need customizable and adaptable discovery policies. In one embodiment, the peer-to-peer platform may be policy-agnostic, and may only provide the basics for discovery. The basics may include one or more core discovery protocols including, but not limited to, a propagate protocol (broadcast within a scope range (subnet or peer group members)), a rendezvous protocol (unicast to a trusted discovery peer) and an invite protocol (reverse discovering).

[0180] A discovery policy may be implemented in a discovery service based on the core discovery protocol. In one embodiment, a discovery service in the core peer-to-peer platform may be used to discover abstractions and/or entities in the peer-to-peer network including, but not limited to, peers, peer groups, peer group policies (group defined services) and pipe endpoints.

[0181] In some embodiments of a peer-to-peer platform, the discovery service may rely on trusted peers (discovery proxies). The discovery service may leverage local neighbors (local propagate). The discovery service may use rendezvous peers (indexes). The discovery service may leave traces in discovery proxies (cache). The discovery service may use net crawling as a last resort (propagate between trusted discovery proxies). In one embodiment, a discovery service may not discover some entities in the peer-to-peer network including, but not limited to, content (large scale; in one embodiment, a content management service may be used for content discovery), metadata (maintain relationship

between data), users, and applications.

[0182] Embodiments of a peer-to-peer platform discovery service may leverage surrounding peers and peer groups, provide meetings points for far away peers and groups, use an asynchronous protocol and provide reverse discovery. The discovery service preferably may be used to find new neighbor peers and provide the ability for a peer to learn about other peer's abilities. Embodiments of a discovery service in the peer-to-peer platform may provide extensibility, spontaneous configuration, adaptive connectivity, a dynamic (i.e. no fixed) network topology, and the ability to reach the "edge of the Internet" (firewall, and NAT).

[0183] Embodiments of a discovery method in the peer-to-peer platform preferably do not require centralized naming (e.g. no DNS). A discovery service preferably may provide predefined meeting points that may be used in platform bootstrapping. The discovery service preferably may support a dynamic environment (peers may come and go). The discovery service preferably may support an unreliable environment (peers may fail). The discovery service preferably may help to adapt to a changing environment through viral behavior. The discovery service preferably may be used to improve performance as a system ages (increase locality). The discovery service preferably may be used in support of security (change of physical location). The discovery service preferably may be used that provides administrationless discovery (zero-admin).

[0184] Embodiments of the peer-to-peer platform discovery service may allow a peer to learn about other peers that discover it. In one embodiment, the peer-to-peer platform discovery service may provide application-managed rendezvous. In one embodiment of the peer-to-peer platform, a peer discovery protocol may support a discovery query message and a discovery response message to be used in the peer discovery process.

[0185] Peer groups need customizable and adaptable discovery policies. One approach to implementing a discovery policy is to start simple and build more complex policies. Embodiments of the peer-to-peer platform discovery service may support discovery methods including, but not limited to:

- Propagate Discovery
 - Unicast to predefined rendezvous
 - Leverage transport dependent multicast (e.g. IP)
- Unicast Discovery
 - Unicast to known rendezvous for forward propagation
 - May be used for reverse Discovery

[0186] The peer-to-peer platform preferably does not mandate exactly how discovery is done. Discovery may be completely decentralized, completely centralized, or a hybrid of the two. Embodiments of the peer-to-peer platform may support discovery mechanisms including, but not limited to:

- LAN-based discovery. This is done via a local broadcast over the subset.
- Discovery through invitation. If a peer receives an invitation (either in-band or out-of-band), the peer information contained in the invitation may be used to discover a (perhaps remote) peer.
- Cascaded discovery. If a peer discovers a second peer, the first peer may, with the permission of the second peer, view the horizon of the second peer to discover new peers, groups, and services.
- Discovery via rendezvous points. A rendezvous point is a special peer that keeps information about the peers it knows about. A peer that can communicate via a rendezvous peer, for example via a peer-to-peer protocol pipe, may learn of the existence of other peers. Rendezvous points may be helpful to an isolated peer by quickly seeding it with lots of information. In one embodiment, a web site or its equivalent may provide information of well-known peer-to-peer protocol rendezvous points.

[0187] In one embodiment, a peer-to-peer platform web of trust may be used. In a web of trust, a peer group creator may select initial discovery proxies, and may delegate to new peer members. Any peer, when trusted, can become a discovery proxy. Discovery proxies may propagate requests between each other for net-crawling discovery. New peers may be untrusted or low-trust peers, and may be typically difficult to find and have limited discovery range (this may help protect against misbehaviors and denial of service attacks). Trusted members are easier to discover. Peers may increase their discovery range as they become more trusted (discovery credential). Some peers may not need to discover beyond their initial net peer group range.

[0188] In one embodiment, a peer may go through a proximity network, which also may be referred to as a subnet or region, to try to find (discover) surrounding peers. The Internet includes the concept of subnets that are physically defined by physical routers that define regions in which computer systems are connected to one another. Within one

of these regions, the peer-to-peer protocol uses multicast or other propagate mechanism to find peers. In one embodiment, a propagate discovery mechanism may be provided where one peer can propagate a discovery request through a local subnet. Peers that are in the subnet may respond to the discovery request. The propagate discovery mechanism may provide primarily close range discovery. In one embodiment, only peers that are in the same physical subnet (region) may respond. "Propagate" is at the conceptual level. Multicast is implemented by TCP/IP to provide propagate capabilities. Other transports may use other methods to implement propagate. For example, Bluetooth provides a different implementation of propagate which is not multicast.

[0189] The core discovery protocol may provide a format for a local peer to send a propagate message (a request to find information about other peers or peer groups in its local region or subnet) and also a format for a response message. A propagate may ask who's there (what peers are in the subnet). One or more peers may decide to respond. Other peers on the subnet may choose not to respond if they don't want to be discovered by the requesting peer. The response message may indicate that a peer is there and that the requesting peer may communicate with it if it wants more information. In one embodiment, the core peer-to-peer platform may define the format of the discovery requests and responses as part of the peer discovery protocol. In one embodiment, the messages may be XML messages.

[0190] One embodiment of a peer-to-peer platform may provide a bootstrapping process for peers. In one embodiment, a new peer may not know any peers or peer groups when bootstrapped. When bootstrapping, the peer may issue a peer discovery propagate message. The new peer is looking for one or more peers in the subnet. The new peer needs to reach some level of connectivity in order to support higher-level operations. From discovered peers, the new peer may acquire information needed to allow the new peer to go further in its bootstrapping process. For example, the new peer may send messages to another peer requesting information on services that the other peer may be aware of that the new peer needs for bootstrapping.

[0191] When the new peer discovers another peer or peers, it may attempt to discover peer groups. This process may be similar to the peer discovery process described above. The new peer may send (e.g. propagate) another discovery message that is configured to discover peer groups. Peers in the proximity network (region) that are aware of a peer group or peer groups may respond to the peer group discovery message, and may return information on the peer group(s) (e.g. peer group advertisements) of which they are aware. The new peer may use this information to determine a peer group or peer groups that it may be interested in joining.

[0192] In one embodiment, a peer group may be configured so that only a subset of peers within a group may have the capabilities to respond to peer group discovery messages and to provide information about the peer group to inquiring peers.

[0193] Peer and peer group discovery may both be implemented by the peer discovery protocol. Peer and peer group discover are more or less at the same level in the P2P platform. In one embodiment, peer discovery may use a message that indicates the discovery is looking for peers, and peer group discovery may use a similar message that indicates the discovery is looking for peer groups.

[0194] In one embodiment, the peer discovery protocol may be required to be implemented in a peer platform, and thus all peers will have the service running. When one peer sends (e.g. propagates) a request, then a receiving peer must send a response, unless it is configured to not respond to at least some requests from at least some peers based upon configuration parameters. In another embodiment, peers may be implemented without the peer discovery protocol. In other words, in this embodiment, peers are not required to implement the peer discovery platform. For example, on some smart devices, peer information and/or peer group information may be preconfigured into the device, and so bootstrapping may be performed on these devices without having to initiate a peer discovery.

[0195] Embodiments of the peer-to-peer platform may implement a discovery mechanism that is more suited for long-range discovery than the propagate method described above. In one embodiment, rendezvous peers may be used in discovery. A rendezvous peer may be described as a meeting point where peers and/or peer groups may register to be discovered, and may also discover other peers and/or peer groups, and retrieve information on discovered peers and/or peer groups. In one embodiment, a peer (any peer) in a peer group may decide to become or may be appointed or elected as a rendezvous peer in the group. The rendezvous peer may be advertised as a meeting point, and may be predefined on peers so that, for example, the peers, when starting up, may know to go to the rendezvous peer to find information about the peer-to-peer network. Rendezvous peers may act as information brokers or centralized discovery points so that peers can find information in an easy and efficient manner. As a peer group grows, a peer may become a rendezvous peer in the group. In one embodiment, a network of rendezvous peers may be constructed that may help to provide long-range discovery capabilities. A rendezvous peer may be aware of at least some of the other rendezvous peers in the network, and a discovery message from a peer may be forwarded from a first rendezvous peer to a second, and so long, to discover peers and/or peer groups that are "distant" on the network from the requesting peer.

[0196] Rendezvous peers may offer to cache advertisements to help others peers discover resources, and may propagate (forward) requests they cannot answer to other known rendezvous peers. Preferably, a rendezvous peer implements at least one of these two functions. The services provided by a rendezvous peer may be different than

message routing. Message routing is performed at a lower level involving multi-hops connections to send a message between any peers in the network. In one embodiment, the forwarding of a request between two rendezvous peers may involve routing to propagate a request between two rendezvous, but this is transparent to the rendezvous service and done underneath.

5 **[0197]** In one embodiment, rendezvous peers may forward requests between each other. A rendezvous may be typically connected to a few other rendezvous peers. There may be as many rendezvous peers as peers in a peer group. Not every peer may be a rendezvous (e.g. if a peer has no caching capabilities or is isolated behind a firewall). In one embodiment, only rendezvous peers may forward a discovery request to another rendezvous peer. This restriction may serve to limit and control the exponential growth of request propagations within the network. Rendezvous
10 peers may thus provide a simple throttle mechanism to control the propagation of requests. In one embodiment, sophisticated rendezvous peers may be deployed to filter and distribute requests for the best usage of network resources.

[0198] In one embodiment, a peer may be pre-configured with a pre-defined set of rendezvous peers. These bootstrapping rendezvous may help the peer discover enough network resources (peers, rendezvous, services) as it needs to support itself. In one embodiment, the pre-configured rendezvous are optional. A peer may be able to bootstrap
15 itself by finding rendezvous or enough network resources in its proximity environment. If a peer does not know the information, it may ask the surrounding peers (hop of 1) if they know the answer. One or more peers may already have the answer. If no surrounding peers know the answer, the peer may ask its rendezvous peers to find advertisements. Peers are recognized as rendezvous peers in their peer advertisements. When a peer discovers a new peer, it can determine if this peer is a rendezvous. A peer may not be required to use all the rendezvous peers that it has discovered.

20 **[0199]** Rendezvous peers may forward requests between themselves. The discovery process continues until one rendezvous peer has the answer or the request dies. There is typically a Time To Live (TTL) associated with the request, so it is not infinitely propagated. As an example, suppose a peer A is attempting to discover a resource R on the network. Peer A issues a discovery request specifying the type (peer, peer group, pipe, service) of advertisements it is looking for. To initiate the Discovery, peer A sends a discovery request message as a single propagate packet to all
25 its available endpoints. The packet may contain the requested peer advertisement, so the receiving peer can respond to the requester. Each discovery request identifies the initiator, and a unique request identification specified by the initiator of the request. When another peer receives the discovery request (assume peer B in this example), if it has the requested R advertisement, it will return to peer A the advertisement for R in a discovery response message. If Peer A does not get response from its surrounding peers (hop of 1), Peer A may send the request to its known rendezvous peers. If the rendezvous peers do not have the advertisement, they can propagate the request to all other rendezvous peers they know. When a rendezvous receives a respond to a request, the rendezvous MAY cache the R advertisement for future usage, before sending it to the requestor.

[0200] In one embodiment, the peer rendezvous capabilities may be embedded in the core discovery protocol of the peer-to-peer platform. Rendezvous peers may be protocol-based, and may broker more information than name servers
35 that typically only broker names of entities. In one embodiment, a rendezvous peer may maintain indexes for entities in the peer-to-peer platform including peers, peer groups, and advertisements. These indexes are dynamic which are created as the peer group community grows and more peers join. As a group joins, some peers may decide to become rendezvous peers to help peers connect with other peers in the group.

40 **[0201]** The rendezvous peer is at the peer level. A rendezvous peer is not a "service". A rendezvous peer may be used as part of an infrastructure to construct services such as a DNS or other centralizing and index services. In one embodiment, services may interact with a rendezvous peer to obtain and/or manipulate information stored on the rendezvous peer to perform some task to make the system act more efficiently.

[0202] In a network of peers, some peers may elect themselves, through the discovery protocol, to become rendezvous peers. A rendezvous peer may act as a broker or discovery message router to route discovery messages to the
45 right place. In other words, a rendezvous may act to route discovery requests to the right rendezvous peers. For example, a rendezvous peer may receive a message requesting information about peers that are interested in baseball. The rendezvous peer may know of another rendezvous peer that specializes in information about baseball. The first rendezvous peer may forward or route the message to the second rendezvous peer. In one embodiment, rendezvous peers may maintain connections to other rendezvous peers in order to provide discovery and routing functionality.

50 **[0203]** Rendezvous peers may support long-range discovery. For example, a first peer is at a remote location from a second peer. For one of these peers to find the other with a mechanism such as web crawling may be time consuming, since there maybe a lot of "hops" between the two peers. Rendezvous peers may provide a shortcut for one of the peers to discover the other. The rendezvous peer, thus, may serve to make the discovery process, in particular long-range discover, more efficient.

55 **[0204]** A peer-to-peer network may be dynamic. Peers and peer groups can come and go. Dynamic identifiers (addresses) may be used. Thus, routes between peers need to be dynamic. Rendezvous peers may provide a method for route discovery between peers that allows routing in the peer-to-peer network to be dynamic. In this method, the rendezvous peers may perform route discovery for peers when the peers send discovery messages to the rendezvous

peers or when a peer is attempting to connect to another peer or peer group that is not in the local region of the peer. This method may be transparent to the requesting peer.

[0205] In one embodiment, the rendezvous peers may be able to cache advertisements. An advertisement may be defined as metadata or descriptions of a resource. An advertisement may include information necessary for an entity to connect to or use the resource, for example a service advertisement may include information for connecting to and using the service. Advertisements may be published to allow other entities to discover them. The rendezvous peer may provide the ability for services and applications to store and cache temporary, e.g. via a lease mechanism, advertisements. This may be used, for example, when one service needs to connect to another service, and needs the pipe endpoint or communication channel that may be used to connect to the service. The pipe endpoint may be included in a service advertisement published on a rendezvous peer. Thus, in one embodiment, the rendezvous peer provides the ability for peers, peer groups, services and applications to advertise pipe endpoints and to discover pipe endpoints of services and applications.

[0206] In one embodiment, the rendezvous protocol may use an index cache (e.g. on a peer serving as a rendezvous proxy). Figure 13 illustrates discovery through a rendezvous peer according to one embodiment. Rendezvous proxy 206 may cache peer 200 and peer group 210 information for peer groups 210A and 210B. Peers 200 in each peer group 210 may then discover each other through rendezvous proxy 206. Rendezvous proxy 206 may itself be a peer and may be a member in one or more peer groups 210. In one embodiment, access to rendezvous proxies 206 may be restricted to peers with rendezvous access privileges. In this embodiment, non-trusted peers (peers without access privileges) may access rendezvous proxies 206 through trusted peers 200 within their peer group 210, or alternatively through other local peers in other peer groups. In one embodiment, the rendezvous protocol may be used across subnets (configurable at the peer group level). In one embodiment, the rendezvous protocol may be used across/through firewalls (e.g. gateways).

[0207] In one embodiment, the peer-to-peer platform may include a propagate policy for use in discovery. Figure 14 illustrates discovery through propagate proxies according to one embodiment. In one embodiment, discovery proxy 208 may control propagation of discovery messages. In Figure 14, discovery proxy 208 may receive discovery messages from peers 200 in peer group 210A and propagate the messages to peers in other groups such as peer group 210B. In one embodiment, access to discovery proxies 208 may be restricted to peers with discovery proxy access privileges. In this embodiment, non-trusted peers (peers without access privileges) may access discovery proxies through trusted peers 200 within their peer group 210, or alternatively through other local peers in other peer groups. In one embodiment, propagation may be controlled using TTL (time to live). In another embodiment, propagation may be controlled using message counts. In one embodiment, the propagate policy may support HTTP gateways (platform configurable). In one embodiment, the propagate policy may be used through firewalls (e.g. need peer activation behind firewalls).

[0208] In one embodiment, the peer-to-peer platform may include an invite policy. In one embodiment, the invite policy may support the adding of new peers and peer groups (e.g. publish advertisements).

[0209] In one embodiment, the peer-to-peer platform may allow the persistent local peer caching of discovery information. In this embodiment, a peer may be allowed to cache advertisements discovered via the peer discovery protocol for later usage. Caching may not be required by the peer-to-peer platform, but caching may be a useful optimization. The caching of advertisements by a peer may help avoid performing a new discovery each time the peer is accessing a network resource. In a highly transient environment, performing the discovery may be necessary. In a static environment, caching may be more efficient.

[0210] In one embodiment, the peer-to-peer platform may support trusted discovery peers. In one embodiment, the peer-to-peer platform may use discovery credentials. In one embodiment, the peer-to-peer platform may allow credential delegation. In one embodiment, the peer-to-peer platform may support propagate proxies. In one embodiment, a propagate proxy may support TTL/message counts. TTL stands for Time To Live (how long the request lives in the system). In one embodiment, a propagate proxy may support net crawling. In one embodiment, a propagate proxy may provide "smart above" routing.

[0211] In one embodiment, a peer preferably does not initiate a new discovery request until the minimum allowable interval between discoveries is reached. This limitation on the maximum rate of discoveries may be similar to the mechanism required by Internet nodes to limit the rate at which ARP requests are sent for any single target IP address. The maximum rate may be defined by each specific implementation transport bindings and exported to the application.

[0212] Figure 15 illustrates using messages to discover advertisements according to one embodiment. A message or messages may be used to get all known, reachable advertisements within a region on the network. This list is preferably not guaranteed to be exhaustive, and may be empty. Named peers may also be located using the peer discovery protocol. A message may include a peer group credential of the probing (requesting) peer that may identify the probing peer to the message recipient. The destination address may be any peer within a region (a propagate message 230) or alternatively a rendezvous peer (a unicast message 232). The response message 234 may return one or more advertisements (e.g. peer advertisements and/or peer group advertisements) that may include "main"

endpoint addresses which may be converted to a string in the standard peer endpoint format (e.g. URI or URL) and also may include a network transport name. It is preferably not guaranteed that a response to a query request will be made. Preferably, the peer discovery protocol does not require a reliable transport. Multiple discovery query requests may be sent. None, one or multiple responses may be received.

[0213] In one embodiment, a discovery query message may be used to send a discovery request to find advertisements (e.g. for peers or peer groups). The discovery query may be sent as a query string (attribute, value) form. A null query string may be sent to match any results. A threshold value may be included to indicate the maximum number of matches requested by a peer. The following is an example of one embodiment of a discovery query message in XML, and is not intended to be limiting:

```
<?xml version="1.0" encoding="UTF-8"?>
<DiscoveryQuery>
  <Credential> Credential </Credential>
  <QueryId> query id</QueryId>
  <Type> request type (e.g. PEER, GROUP, PIPE, SERVICE, CONTENT) </Type>
  <Threshold> requested number of responses </Threshold>
  <PeerAdv> peer advertisement of requestor </PeerAdv>
  <Attr> attribute </Attr>
  <Value> value </Value>
</DiscoveryQuery>
```

[0214] Embodiments of a discovery query message may include, but are not limited to, the following fields:

- Credential: The credential of the sender
- QueryId: Query identifier
- Type: specifies which advertisements are returned
- Threshold: requested number of responses
- PeerAdv: peer advertisement of requestor
- Attr: specifies the query attribute
- Value: specifies the query value

[0215] In one embodiment, the value tag is only present if the Attr tag field is present. Both the Attr and Value tag may be omitted.

[0216] In one embodiment, a discovery response message may be used to send a discovery response message to answer a discovery query message. The following is an example of one embodiment of a discovery response message in XML, and is not intended to be limiting:

```

<?xml version="1.0" encoding="UTF-8"?>
<DiscoveryResponse>
  <Credential> Credential </Credential>
  <QueryId> query id</QueryId>
  <Type> request type (e.g. PEER, GROUP, PIPE, SERVICE, CONTENT) </Type>
  <Attr> Attribute </Attr>
  <Value> value </Value>
  <Responses>
    (peer, peer group, pipe, service or content advertisement response)
  </Responses>
  <Responses>
    (peer, peer group, pipe, service or content advertisement response)
  </Responses>
</DiscoveryResponse>

```

[0217] Embodiments of a discovery response message may include, but are not limited to, the following fields:

- Credential: The credential of the sender
- QueryId: Query identifier
- Type: specifies which advertisements are returned
- Attr: specifies the query attribute
- Value: specifies the query value
- Responses: advertisement responses. The advertisement may be a peer, peer group, pipe, content or service advertisement.

[0218] In one embodiment, the value tag is only present if the Attr tag field is present. Both the Attr and Value tag may be omitted.

[0219] In one embodiment, if an XML advertisement document is embedded into another XML document, the XML document separators must be dealt with. This may be done using the standard XML escaping rules. For example, '<' becomes '<'; '>' becomes '>'; and '&' becomes '&'.

Reverse discovery

[0220] Reverse discovery means that, in a peer-to-peer network, when a first entity (e.g. a peer) discovers a second entity (e.g. another peer), the second entity may also discover the first entity from the discovery initiated by the first entity. This may also be referred to as "mutual discovery". In most traditional systems, discovery is typically one-directional. In the peer-to-peer world, reverse discovery is important because, by definition, all "peers" are equal (i.e. it is typically not a hierarchical system). In one embodiment, there may be different levels of discovery for peers. For example, a peer may be configured to remain anonymous when discovering other peers or to always support reverse discovery. In one embodiment, a peer initiating a discovery may also be configured to deny discovery to another peer if the other peer is configured or chooses to remain anonymous. In one embodiment, a peer may also be configured to or may choose to deny discovery by other peers that wish to remain anonymous.

Invitations

[0221] One embodiment of the discovery protocol may also provide methods by which a peer can "advertise" itself, for example when joining a peer-to-peer network. For example, a peer may send an email message, by telephone, by "traditional" mail, or by other methods to other peers it discovers or is preconfigured to know about to advertise its presence and willingness to be contacted by other peers. This is done outside of the discovery method, and may be performed by any external medium. A peer who receives an invitation from a peer may have a capability to add or enter the new peer to a list or database of peers that it knows about. When the peer later restarts, these peers may

be among the preconfigured peers that the peer knows about. In one embodiment, a peer may have a "notify" or "invitation" interface to allow a user to initiate invitations. In one embodiment, the peer-to-peer platform may provide import and export capabilities for invitations. In one embodiment, the invitations may be implemented as documents external to the peer-to-peer system that may be exported from one peer and imported into another peer. In one embodiment, the invitations may be in a format that enables the exporting and importing. In one embodiment, the invitations may be in XML format. In one embodiment, an interface may be provided to allow the manual entering of invitation information. Importing the invitation may create a peer-to-peer platform document that may then be used by the peer. The format of exported documents may depend on the platform on which the peer is implemented.

Peer Resolver Protocol

[0222] In one embodiment, the peer-to-peer platform may include a peer resolver protocol that may allow a peer to send preferably simple, generic search queries to one or more peer services. In one embodiment, only those peers that have access to data repositories and that offer advanced search capabilities typically implement this protocol. Each service may register a handler in the peer group resolver service to process resolver query requests. Resolver queries may be demultiplexed to each service. Each service may respond to a peer via a resolver response message. It is important to point the differences between the peer discovery protocol and the peer resolver protocol. The peer discovery protocol is used to search for advertisements to bootstrap a peer, and discover new network resources. The peer resolver protocol is a generic service that services query protocols. The peer resolver protocol may be used by a service on a peer to interact with a service on another peer.

[0223] The peer resolver protocol may enable each peer to send and receive generic queries to find or search for peer, peer group, pipe or service specific information such as the state of a service or the state of a pipe endpoint. Preferably, each resolver query has a unique service handler name to specify the receiving service, and a query string to be resolved by the service. The peer resolver protocol preferably provides a generic mechanism for peers to send queries and receive responses. The peer resolver protocol preferably removes the burden for registered message handlers by each service and set message tags to ensure uniqueness of tags. The peer resolver protocol preferably ensures that messages are sent to correct addresses and peer groups. The peer resolver protocol preferably performs authentication and verification of credentials and the dropping of rogue messages. Preferably, there is no guarantee that a response to a resolver query request will be made. Preferably, a peer is not required to respond to a resolver query request. Preferably, a reliable transport is not required by the peer resolver protocol. In one embodiment, multiple resolver query messages may be sent. None, one or multiple responses may be received.

[0224] In one embodiment, propagating a query to the next set of peers may be delegated to the peer rendezvous protocol. The rendezvous service may be responsible for determining the set of peers that may receive a message being propagated, but may not re-propagate an incoming propagated message. The decision of propagating a message one step further may be left to the service handling the message. The peer rendezvous protocol's policy may be that if the query handler does not instruct the peer rendezvous protocol to discard the query, and if the local peer is a rendezvous, then the query is re-propagated (within the limits of loop and time-to-live rules that may be enforced by the rendezvous service). In addition, if instructed by the query handler, an identical query may be issued with the local peer as the originator.

[0225] Figure 16 illustrates one embodiment of using peer resolver protocol messages between a requesting peer 200A and a responding peer 200B. In one embodiment, a resolver query message 236 may be used to send (unicast) a resolver query request to a service on another member 200B of a peer group. In one embodiment, the resolver query may be sent as a query string to a specific service handler. Preferably, each query has a unique identifier. The query string may be any string that may be interpreted by the targeted service handler. A resolver response message 238 may be sent (unicast) to the requesting peer 200A by the service handler. The following is an example of one embodiment of a resolver query message in XML, and is not intended to be limiting:

```
<?xml version="1.0" encoding="UTF-8"?>
<ResolverQuery>
  <Credential> Credential </Credential>
  <HandlerName> name of handler </HandlerName>
  <QueryId> incremental query Id </QueryId>
  <Query> query string </Query>
</ResolverQuery>
```

[0226] Embodiments of a resolver query message may include, but are not limited to, the following fields:

- Credential: The credential of the sender
- QueryId: Query identifier
- 5 • Handler Name: service the query needs to be passed
- Query: query string

[0227] A resolver response message may be returned in response to a resolver query message. The following is an example of one embodiment of a resolver response message in XML, and is not intended to be limiting:

```

10      <?xml version="1.0" encoding="UTF-8"?>
      <ResolverResponse>
15          <Credential> Credential </Credential>
          <HandlerName> name of handler </HandlerName>
          <QueryId> query Id </QueryId>
          <Response> response </Response>
20      </ResolverResponse>

```

[0228] Embodiments of a resolver response message may include, but are not limited to, the following fields:

- 25 • Credential: The credential of the sender
- QueryId: Query Id (long as a String)
- HandlerName: service the query needs to be passed
- Response: response String

30 Peer Information Protocol

[0229] Once a peer is located, its capabilities and status may be of interest. In one embodiment, the peer-to-peer platform may include a peer information protocol that may allow a peer to learn about other peers' capabilities and status. For example, a peer can send a ping message to see if another peer is alive. A peer may also query another peer's properties where each property has a name and a value string. Preferably, a peer is not required to respond to a peer information protocol request.

[0230] Figure 17 illustrates one embodiment of using peer information protocol messages between a requesting peer 200A and a responding peer 200B. In one embodiment, to see if peer 200B is alive (i.e. responding to messages), peer 200A may be sent a ping message 240. The ping message 240 may include a destination address that is peer 200B's "main" endpoint returned during discovery, for example. The message may also include a group membership credential of the requesting peer 200A that may identify the probing peer 200A to the message recipient 200B. The message may also contain an identifier unique to the sender. This identifier is preferably returned in the response message 242. Response message 242 may include information about peer 200B, including information on the status of the peer 200B. If peer 200B responds with a message 242, this may indicate to peer 200A that peer 200B is "alive" and thus currently responding to messages.

[0231] In one embodiment, messages may be used to get a list of named control "properties" exported by a peer. A property is a "knob" used to get information or configuration parameters from the peer. All properties are preferably named (by a string), and are preferably "read-only". In one embodiment, higher-level services may offer "read-write" capability to the same information, given proper security credentials. Each property preferably has a name and a value string. Read-write widgets may allow the string value to be changed, while read-only widgets do not. In one embodiment, the peer information protocol only gives read access. The destination address is a peer's main endpoint that may have been returned in a discovery response message.

[0232] Preferably, a reliable transport is not required by the peer information protocol. In one embodiment, multiple peer information messages may be sent. None, one or multiple responses may be received.

55 [0233] In one embodiment, a ping message may be sent to a peer to check if the peer is alive and/or to get information about the peer. The ping option may define the response type returned. In one embodiment, a full response (peer advertisement) or a simple acknowledge response (alive and uptime) may be returned. The following is an example of one embodiment of a ping message in XML, and is not intended to be limiting:


```

<?xml version="1.0" encoding="UTF-8"?>
<Ping>
  <Credential> Credential </Credential>
  <SourcePid> Source Peer Id </SourcePid>
  <TargetPid> Target Peer Id </TargetPid>
  <Option> type of ping requested</Option>
</Ping>

```

[0234] In one embodiment, a peer information response message may be used to send a response message in response to a ping message. The following is an example of one embodiment of a peer information response message in XML, and is not intended to be limiting:

```

<?xml version="1.0" encoding="UTF-8"?>
<PeerInfo>
  <Credential> Credential </Credential>
  <SourcePid> Source Peer Id </SourcePid>
  <TargetPid> Target Peer Id </TargetPid>
  <Uptime> uptime</Uptime>
  <TimeStamp> timestamp </TimeStamp>
  <PeerAdv> Peer Advertisement </PeerAdv>
</PeerInfo>

```

Peer Membership Protocol

[0235] In one embodiment, the peer-to-peer platform may include a peer membership protocol that may allow a peer to join or leave peer groups, and to manage membership configurations, rights and responsibilities. This protocol may allow a peer to obtain group membership requirements (such as an understanding of the necessary credential for a successful application to join the group), to apply for membership and receive a membership credential along with a full group advertisement, to update an existing membership or application credential, and to cancel a membership or an application credential. In one embodiment, authenticators and/or security credentials may be used to provide the desired level of protection.

[0236] In one embodiment, the process of joining a peer group may include obtaining a credential that is used to become a group member. In one embodiment, the process of joining a peer group may include obtaining a "form" listing the set of requirements asked of all group members. In one embodiment, this form may be a structured document (e.g. a peer group advertisement) that lists the peer group membership service.

[0237] In one embodiment, the peer membership protocol may define messages including, but not limited to, an apply message, a join message, an acknowledgement (ACK) message, a renew message, and a cancel message. A peer membership protocol apply message may be sent by a potential new group member to the group membership application authenticator. The authenticator's endpoint is preferably listed in the peer group advertisement of every member. In one embodiment, a successful response from the group's authenticator may include an application credential and a group advertisement that preferably lists, at a minimum, the group's membership service. In one embodiment, the apply message may include, but is not limited to, the current credential of the candidate group member and the peer endpoint for the peer group membership authenticator to respond to with an acknowledgement (ACK) message.

[0238] The following is an example of one embodiment of a peer membership protocol apply message in XML, and is not intended to be limiting:

```

<?xml version="1.0" encoding="UTF-8"?>
<MembershipApply>
  <Credential> Credential of requestor </Credential>
  <SourcePid> Source pipe identifier </SourcePid>

  <Authenticator> Authenticator pipe advertisement </Authenticator>
</MembershipApply>

```

[0239] A peer membership protocol join message may be sent by a peer to the peer group membership authenticator to join a group. The peer preferably passes an application credential (from an apply response ACK message) for authentication purposes. A successful response from the group's authenticator preferably includes a full membership credential and a full group advertisement that lists, at a minimum, the group's membership configurations requested of full members in good standing. The message may include a credential (application credential of the applying peer: see ACK message). This credential may be used as the application form when joining. The message may also include the peer endpoint for the authenticator to respond to with an ACK message.

[0240] The following is an example of one embodiment of a peer membership protocol join message in XML, and is not intended to be limiting:

```

<?xml version="1.0" encoding="UTF-8"?>
<MembershipJoin>
  <Credential> Credential of requestor </Credential>
  <SourcePid> Source pipe identifier </SourcePid>
  <Memberssship> membership pipe advertisement </Membership>
  <Identity> identity </Identity>
</MembershipJoin>

```

[0241] A peer membership protocol ACK message is an acknowledge message that may be used for both join and apply operations. A peer membership protocol ACK message may be sent back by the membership authenticator to indicate whether or nor the peer was granted application rights to the peer group if the peer is applying, or full membership to the peer group if peer is attempting to join. In one embodiment, an ACK message may also be sent in response to peer membership protocol renew messages and cancel messages. The message may include a credential (an application or membership credential allocated to the peer by the peer group authenticator). The message may also include a more complete peer group advertisement that may provide access to further configurations. In one embodiment, not all configuration protocols are visible until the peer has been granted membership or application rights. Some configurations may need to be protected. Also, depending on the peer credential, the peer may not have access to all the configurations.

[0242] The following is an example of one embodiment of a peer membership protocol ack message in XML, and is not intended to be limiting:

```

<?xml version="1.0" encoding="UTF-8"?>
<MembershipAck>
  <Credential> Credential </Credential>
  <SourcePid> Source pipe identifier </SourcePid>
  <Memberssship> membership pipe advertisement </Membership>
  <PeerGroupAdv> peer group advertisement </PeerGroupAdv>
  <PeerGroupCredential> credential granted </PeerGroupCredential>
</MembershipAck>

```

[0243] A peer membership protocol renew message may be sent by a peer to renew its credential (membership or application) access to the peer group. An ACK (acknowledgement) message may be returned with a new credential and lease if the new is accepted. The renew message may include, but is not limited to, a credential (a membership or application credential of the peer) and the peer endpoint to which an ACK response message may be sent.

[0244] The following is an example of one embodiment of a peer membership protocol renew message in XML, and is not intended to be limiting:

```
<?xml version="1.0" encoding="UTF-8"?>
<MembershipRenew>
  <Credential> Credential </Credential>
  <SourcePid> Source pipe identifier </SourcePid>
  <Memberssship> membership pipe advertisement </Membership>
</MembershipRenew>
```

[0245] A peer membership protocol cancel message may be sent by a peer to cancel the peer's membership or application rights in a peer group. The message may include, but is not limited to, a credential (a membership or application credential of the peer) and the peer endpoint to send an ACK message. In one embodiment, an ACK to a cancel may include a response status indicating the cancel was accepted.

[0246] The following is an example of one embodiment of a peer membership protocol cancel message in XML, and is not intended to be limiting:

```
<?xml version="1.0" encoding="UTF-8"?>
<MembershipCancel>
  <Credential> Credential </Credential>
  <SourcePid> Source pipe identifier </SourcePid>
  <Memberssship> membership pipe advertisement </Membership>
</MembershipCancel>
```

Pipe Binding Protocol

[0247] In one embodiment, the peer-to-peer platform may include a pipe binding protocol that may allow a peer to find the physical location of a pipe endpoint and to bind a pipe advertisement to the pipe endpoint, thus indicating where messages actually go over the pipe. A pipe is conceptually a virtual channel between two pipe endpoints (input and output pipes) and may serve as a virtual link between two or more peer software components (e.g. services or applications).

[0248] A pipe may be viewed as an abstract, named message queue that supports a number of abstract operations such as create, open, close, delete, send, and receive. The pipe virtual link (pathway) may be layered upon any number of physical network transport links such as TCP/IP. Each end of the pipe may work to maintain the virtual link and to reestablish it, if necessary, by binding endpoints or finding the pipe's currently bound endpoints.

[0249] Actual pipe implementations may differ, but peer-to-peer platform-compliant implementations preferably use the pipe binding protocol to bind pipes to pipe endpoints. In one embodiment, during the abstract create operation, a local peer binds a pipe endpoint to a pipe transport. In another embodiment, bind may occur during the open operation. Unbind occurs during the close operation. In one embodiment, each peer that "opens" a group pipe may make an endpoint available (binds) to the pipe's transport. Messages are preferably only sent to one or more endpoints bound to the pipe. Peer members that have not opened the pipe preferably do not receive or send any messages on that pipe. In one embodiment, when some peer software wants to accept incoming pipe messages, the receive operation may remove a single message in the order it was received, not in the order it was sent. In one embodiment, a peek operation may be used as a mechanism to see if any message(s) has arrived in the pipe's queue.

[0250] In one embodiment, the pipe binding protocol may define messages including, but not limited to, a query message and a response message. In one embodiment, a pipe binding protocol query message may be sent by a peer pipe endpoint to find a pipe endpoint bound to the same pipe advertisement. The following is an example of one embodiment of a pipe binding protocol query message in XML, and is not intended to be limiting:

<?xml version="1.0" encoding="UTF-8"?>

<PipeBindingQuery>

<Credential> query credential </Credential>

<Peer> optional tag. If present, it may include the peer identifier of the only peer that should answer the request.

</Peer>

<Cached> true if the reply can come from a cache </Cached>

<PipeId> pipe identifier to be resolved </PipeId>

</PipeBindingQuery>

[0251] In one embodiment, the requestor may ask that the information not be obtained from a cache. This is to obtain the most up-to-date information from a peer to address stale connection. The Peer field specifies a peer identifier. This peer is the one that should respond to the query. There is preferably no guarantee that a response to a pipe binding request will be made. Preferably, a peer is not required to respond to a binding request. Preferably, a reliable transport is not required. In one embodiment, multiple binding query messages may be sent. None, one or multiple responses may be received.

[0252] In one embodiment, a pipe binding protocol response message may be sent to the requesting peer by each peer bound to the pipe in response to a query message. The following is an example of one embodiment of a pipe binding protocol response message in XML, and is not intended to be limiting:

<?xml version="1.0" encoding="UTF-8"?>

<PipeBindingAnswer>

<Credential> credential </Credential>

<PipeId> pipe id resolved </PipeId>

<Peer> peer URI where a corresponding InputPipe has been created </Peer>

<Found> true: the InputPipe does exist on the specified peer (ACK)

false: the InputPipe does not exist on the specified peer (NACK)

</Found>

</PipeBindingAnswer>

Endpoint Routing Protocol

[0253] In one embodiment, the peer-to-peer platform may include an endpoint routing protocol. The endpoint routing protocol may be used by peers to send messages to router peers requesting available routes for sending message(s) to destination peers.

[0254] A peer-to-peer platform network is typically an ad hoc, multi-hops, and adaptive network by nature. Connections in the network may be transient, and message routing may be nondeterministic. Routes may be unidirectional and change rapidly. Peers may appear and leave frequently. Two communicating peers may not be directly connected to each other. Two communicating peers may need to use router peers to route messages depending on the network topology. For example, the two peers may be on different network transports, or the peers may be separated by a firewall or a NAT (Network Address Translation) router. A peer behind a firewall may send a message directly to a peer outside a firewall. But a peer outside the firewall cannot establish a connection directly with a peer behind the firewall.

[0255] The endpoint routing protocol may define a set of request/query messages that is processed by a routing service to help a peer route messages to its destination. When a peer is asked to send a message to a given peer endpoint address, it may look in its local cache to determine if it has a cached route to this peer. If the peer does not find a route, it may send a route resolver query message to available peer routers requesting route information. A peer may have access to as many peer routers as it can find, or optionally a peer may be pre-configured to access certain routers.

[0256] Peer routers provide the low-level infrastructures to route a message between two peers in the network. Any number of peers in a peer group may elect themselves to become peer routers for other peers. Peers routers offer the

ability to cache route information, as well as bridging different physical (different transport) or logical (firewall and NAT) networks. A peer may dynamically find a router peer via a qualified discovery search. A peer may find out if a peer it has discovered is a peer router via the peer advertisement properties tag.

[0257] When a peer router receives a route query, if it knows the destination (a route to the destination), it may answer the query by returning the route information as an enumeration of hops. The message may be sent to the first router and that router may use the route information to route the message to the destination peer. The route may be ordered from the next hop to the final destination peer. At any point the routing information may be obsoleted, requiring the current router to find a new route.

[0258] The peer endpoint may add extra routing information to the messages sent by a peer. When a message goes through a peer, the endpoint of that peer may leave its trace on the message. The trace may be used for loop detection and to discard recurrent messages. The trace may also be used to record new route information by peer routers.

[0259] In one embodiment, the endpoint routing protocol may provide the last resort routing for a peer. More intelligent routing may be implemented by more sophisticated routing services in place of the core routing service. High-level routing services may manage and optimize routes more efficiently than the core service. In one embodiment, the hooks necessary for user defined routing services to manipulate and update the route table information (route advertisements) used by the peer router may be provided by the endpoint routing protocol. Thus, preferably complex route analysis and discovery may be performed above the core by high-level routing services, and those routing services may provide intelligent hints to the peer router to route messages.

[0260] Router peers may cache route information. Router peers may respond to queries with available route information. Route information may include a list of gateways along the route. In one embodiment, any peer may become a router peer by implementing the endpoint routing protocol. The following is an example of one embodiment of route information in XML, and is not intended to be limiting:

```
<?xml version="1.0" encoding="UTF-8"?>
<EndpointRouter>
  <Credential> credential </Credential>

  <Src> peer identifier of the source </Src>
  <Dest> peer identifier of the destination </Dest>
  <TTL> time to live </TTL>
  <Gateway> ordered sequence of gateway </Gateway>
  .....
  <Gateway> ordered sequence of gateway </Gateway>
</EndpointRouter>
```

[0261] The time-to-live parameter specifies how long this route is valid. The creator of the route can decide how long this route will be valid. The gateways may be defined as an ordered sequence of peer identifiers that define the route from the source peer to the destination peer. The sequence may not be complete, but preferably at least the first gateway is present. The first gateway is sufficient to initially route the messages. The remaining gateway sequence is preferably optional.

[0262] The endpoint routing protocol may provide messages including, but not limited to, a route request message and a route answer message from the router peer. In one embodiment, a peer may send a route request message to a router peer to request route information. Route information may be cached or not cached. In some cases, the route query request message may indicate to bypass the cache content and thus to search dynamically for a route. Preferably, it is not guaranteed that a route response will be received after a query is sent. The following is an example of one embodiment of a route query request message in XML, and is not intended to be limiting:

```

<?xml version="1.0" encoding="UTF-8"?>
<EndpointRouterQuery>
  <Credential> credential </Credential>
  <Dest> peer identifier of the destination </Dest>
  <Cached> true: if the reply can be a cached reply
             false: if the reply must not come from a cache
  </Cached>
</EndpointRouterQuery>

```

[0263] In one embodiment, a router peer may send a route answer message to a peer in response to a route information request. The following is an example of one embodiment of a route answer message in XML, and is not intended to be limiting:

```

<?xml version="1.0" encoding="UTF-8"?>
<EndpointRouterAnswer>
  <Credential> credential </Credential>
  <Dest> peer id of the destination </Dest>
  <RoutingPeer> Peer identifier of the router that knows a route to DestPeer </RoutingPeer>
  <RoutingPeerAdv> Advertisement of the routing peer </RoutingPeerAdv>
  <Gateway> ordered sequence of gateways </Gateway>
  .....
  <Gateway> ordered sequence of gateways </Gateway>
</EndpointRouterAnswer>

```

Routing

[0264] A peer-to-peer platform preferably provides a mechanism or mechanisms for searching and accessing peers, peer groups, content, services and other information in a dynamic topology of peers and peer groups, where peers and peer groups can come and go. In one embodiment, peers and peer groups may come and go potentially with limited or no control and notification. Peers may connect to a peer-to-peer network through various wired and wireless protocols, including "not connected" protocols such as may be used by mobile consumer devices such as pagers and PDAs. Peers may also have to cross boundaries, for example boundaries created by firewalls and NAT (Network Address Translation) routers, to connect to other peers.

[0265] An application that supports the peer-to-peer platform preferably is able to help in routing and discovering. Some of the information needed to accomplish routing and discovering may be only known by the application. For example, the application may support a special type of data as content, and so the application may best "know" how to discover items of this special content. Also, the application may have a better knowledge of the topology (related to the nature of the application and/or peer group) than the core peer-to-peer platform.

[0266] In one embodiment, in order to bootstrap the system, and also in order to have a fallback mechanism if an application cannot or does not support one or more of the tasks, the core peer-to-peer protocols may provide a discovery and router mechanism for discovering peers and other core abstractions such as advertisements, pipes, and peer groups. The discovery and routing mechanism of the peer-to-peer platform preferably uses as few protocols as possible, is simple, and makes use of underlying optimizations when available. Hooks into the core discovery and router mechanism may be provided so that applications and services may participate in the discovery and router mechanisms, for example, by passing information to the core discovery and router mechanism. In one embodiment, an application or service may be allowed to override the core discovery and router mechanism with its own custom mechanism.

[0267] In one embodiment, the core discovery and router mechanism may be based on web crawling. Web crawling may be well suited for use in self-organizing networks such as peer-to-peer networks. In one embodiment, peers may be configured to participate or not to participate in the discovery and router mechanism, and may be configured as to

the level of involvement in the process. In one embodiment, a peer may decide whether to participate in a discovery or routing task depending on the peer's configuration in the peer-to-peer network. In one embodiment, the configuration may be determined using an automated detection of the position of the peer on the network and a network configuration wizard tool.

[0268] Web crawling may not create bottlenecks such as may be created by the mechanism of a client knowing a server and always going to the same server to find and retrieve information (e.g. DNS, NFS etc.). Even if a server is replicated, like DNS, it is still a centralized server. If all the known instances of the server are not reachable, a client may lose access to the server, even if another (but unknown) server is, indeed, available. In a point-to-point network, the information a peer is looking for is generally "close by" or may eventually be "close by", so web crawling may not go too far.

[0269] Figure 18 illustrates several core components and how they interact for discovery and routing according to one embodiment. Application 300 may use discovery 308 to find peers, peer groups advertisements, and other entities on the peer-to-peer network, and may also publish pipe, peer, peer group, service, and other advertisements for access by other peers, applications and services on the peer-to-peer network. In one embodiment, the endpoint 310 may be responsible for exchanging messages between peers that are directly "connected" to each other (i.e. the peers can reach each other without any routing and/or discovering). When available, multicast may be used to discover peers that the endpoint can reach (multicast is a mechanism which has been introduced in IP in order to optimize this kind of process). In addition to that, or when multicast is not available, A rendezvous and invitation mechanism may also be provided. The rendezvous and invitation method may be used, for example, if multicast is not available. For example, HTTP does not provide multicast capabilities.

[0270] The endpoint router 312 may manage a cache of routes, for example routes to remote peers. In one embodiment, the endpoint router 312 may be configured from caching no routes to caching all routes it is aware of, depending on what the configuration wizard has decided with user control. The endpoint router 312 may also forward (route) messages depending on what is found in the cache, and what has been configured. For instance, the endpoint router 312 may be configured to route search (propagate) requests or to not route the requests.

[0271] In one embodiment, the generic resolver 308 is a protocol that implements a sort of RPC (query/response) protocol on top of the endpoint 310. Discovery 306 and pipe resolver 304 may use the generic resolver. In one embodiment, discovery 306 may be responsible for searching, caching and generating core advertisements (e.g. peer, peer group, and pipe advertisements). Discovery 306 may use the generic resolver 308 to send query messages and to receive answers. In one embodiment, discovery 306 may be aware of rendezvous peers and may have an invitation mechanism that may be used to assist the generic resolver 308. In one embodiment, the pipe resolver 304 may be responsible for localizing the receiving end of a pipe 302 given a pipe advertisement. In one embodiment, the pipe resolver 304 does not search for a pipe advertisement. In one embodiment, the pipe resolver 304 may be configured to manage a cache of the locations of the receiving ends (i.e. receiving peers) of the pipe 302.

[0272] The pipe protocol may use the endpoint 310 for transferring messages (with the potential help of the endpoint router 312) between the sending end of the pipe 302, and the receiving end of the pipe 302. In one embodiment, a pipe 302 may be viewed as an endpoint 310 that has not been bound to a particular peer. In one embodiment, a pipe 302 may be moved seamlessly from one peer to another. In one embodiment, a pipe 302 may also provides uniqueness that may not be provided by an endpoint 310 since a pipe identifier is unique in time and space, and an endpoint 310, being a network address, may not be.

[0273] A discovery and router mechanism based on web crawling may be time-expensive, and higher level protocols (such as applications) may have information that the core is not aware of that may help in the web crawling process. In one embodiment, to enable applications to participate in the process, components of the core mechanism may provide hooks that enable the applications to assist in the process (e.g. by providing information). Some transport protocols such as HTTP may be configured for and/or dynamically learn about web rendezvous peers it can use. An application may be provided access to the list of rendezvous peers. In one embodiment, an application may be allowed to set/unset routes in an endpoint router 312. Each route may be qualified to route or not route propagate messages such as web crawling messages and/ or unicast messages. The endpoint router 312 may be viewed as a route cache manager, which is may be controlled by an endpoint 310 and/or other entities that may need to control it. In one embodiment, an endpoint router 312 may be able to discover unknown routes from applications. In one embodiment, discovery 308 may be configured (statically and/or dynamically) to control the nature and the amount of data that it manages. In one embodiment, discovery 308 may be taught where to go search, or where not to go search. In one embodiment, discovery 308 may make an "upcall" to a search/retrieve mechanism. In one embodiment, a pipe resolver 304 may manage a cache of input pipes (receiving ends). In one embodiment, pipe resolver 304 may be accessed by applications to set/unset entries in the cache.

Router peers

[0274] Figure 19 illustrates one embodiment of message routing in a peer-to-peer network that uses the peer-to-peer platform. Peers 200 in peer groups 210A and 210B may communicate with each other through one or more router peers 244. In one embodiment, message routing may route messages to "unreachable" peers, i.e. may allow messages sent from a peer 200 to reach peers 200 that are otherwise unreachable. Networks may be partitioned by firewalls, NAT (Network Address Translation) routers, etc. Message routing may allow messages to be delivered in partitioned networks. Message routing may also allow peers 200 separated by one or more partitions to participate in the same peer group(s) 210. Message routing preferably provides optimized message delivery, for example by optimizing routes between peers 200. Message routing preferably allows for an adaptive peer-to-peer network (e.g. peers may move to remote locations and still receive messages). Message routing preferably provides load balancing. In one embodiment, any peer may be a router peer 244.

[0275] One embodiment may provide for HTTP routing servers. In one embodiment, HTTP routers may provide for message routes that traverse firewalls. In one embodiment, HTTP routers may provide NAT support. In one embodiment, HTTP routers may act as message gateways (TTL). TTL stands for Time To Live (how long the request lives in the system).

[0276] The widespread use of NAT (Network Address Translation) and firewalls may affect the operation of many P2P systems. It also may affect the peer-to-peer platform. In particular, a peer outside a firewall or a NAT gateway cannot discover peers inside the firewall or the NAT gateway. In the absence of getting system administrators to let the peer-to-peer platform traffic through (say by opening a special incoming port at the firewall or gateway), possible methods to deal with this problem include, but are not limited to:

- In one embodiment, peers inside firewalls may be asked to initiate connections to peers outside the firewall.
- In one embodiment, peer nodes may be set up that operate like mailbox offices where traffic to a peer inside the firewall is queued up to be picked up at a designated relay peer outside the firewall. The peer inside the firewall can initially reach outside the firewall, select a relay peer, and widely advertise this fact.

[0277] Later, it can periodically contact the relay peer to retrieve messages.

[0278] One embodiment of the peer-to-peer platform may provide router peers. The router peers may be at a lower level than rendezvous peers. The router peers may provide "pure" message routing. By looking at the destination and source addresses, the router peer may determine where a message needs to be sent. In one embodiment, a router peer may call or access a rendezvous peer to "discover" information about peers, etc. In other words, the router peer may access information from a rendezvous peer to use the information in routing messages.

[0279] In one embodiment, router peers may provide the lowest message routing layer in the peer-to-peer platform. Routing may involve complex topologies. For example, the routing peers may provide a method to route across a firewall, particularly from peers outside the firewall to peers inside the firewall. A peer cannot send a message directly to another peer behind a firewall, since by definition there may be no direct route from a peer outside the firewall to a peer inside the firewall. A router peer may route messages to a gateway peer (a mailbox server where messages for peers behind the firewall may be temporarily stored). In one embodiment, the gateway peer may be a router peer acting as a gateway. The peers behind the firewall may periodically poll the mailboxes provided by the gateway peer to determine if someone has tried to contact them (i.e. are there any messages in my mailbox?). Note that a "pipe" provides an abstraction at a higher level than the message routing provided by router peers, and thus, a pipe may be an abstraction across the network topology between peers, for example peers on opposite sides of a firewall, through which the peers may communicate. At the lowest level, one or more router peers may discover and establish the actual communications route between the peers. This level, however, may be transparent to the peers, who only "see" the pipes.

[0280] In one embodiment, a router peer may build a route table. The router peer may keep information about routes that it discovers and store them in the route table. This allows the router peer to build a knowledge base (the route table) about the network topology as more messages flow on the system. This information may be used by the router peer to discover and establish optimal routes between entities in the network, and may increase its ability to reach other peers.

[0281] A router peer may access another router peer it is aware of to get route information. The route information may be described as a stacked set of destinations (and the routes to the destinations). In one embodiment, the information the router peer stores on a particular route may be incomplete, because the router peer may only know about the route up to a certain point. For example, the router peer may know about a first portion of a route up to another router peer, which knows about the next portion of the route, and so on.

[0282] In one embodiment, each peer has a unique peer ID that is independent of, and is not assigned to, fixed addresses. Peers may move around. Therefore, the peer-to-peer network topology may be dynamic, and may change

every time a peer goes away or moves. Thus, the routing method provided by the router peers is preferably dynamic to support the dynamic topology. When a peer moves and reconnects, the peer is recognized as the same peer that was previously connected elsewhere in the network. This process may use the unique ID of the peer to indicate that the peer is the same one that was previously connected elsewhere. In one example, when a peer moves, it may go through a discovery process to discover peers and rendezvous peers in its new local subnet or region. If the peer wishes to join a peer group that it used at its previous location, it may then attempt to discover other peers that have knowledge of the peer group or other peers in the peer group. The message may be passed through several router peers until it may reach a router peer that has knowledge about the peer group (e.g. a route to the peer group) to return to the requesting peer. For example, a user with a laptop may fly from a home office to another city. When the user connects to the network in the other city, a route may be established, through the services provided by router peers, to the home office network peer group. The user may then access email and other services provided by the peer group. From the user's standpoint, this process may seem automatic. For example, the user may not be required to "dial in" or connect remotely to an ISP to access the office as is required in typical networks using static addressing.

[0283] In one embodiment, when a peer becomes a router peer, it may access a stored route table as a starting point. In one embodiment, the peer may start from scratch with an empty route table. In one embodiment, the peer, when it becomes a router peer, may initiate a discovery of other router peers and/or rendezvous peers to get as much connectivity information to key peers in the network as possible.

[0284] In one embodiment, every peer may have knowledge of at least one router peer. In one embodiment, there may be a "universal router" that many or all peers may be aware of that may be accessed when a peer cannot find anyone. The universal router may be able to put the peer into contact with somebody (e.g. another peer) to help in the bootstrapping process.

Security

[0285] The security requirements of a P2P system are very similar to any other computer system. The three dominant requirements are confidentiality, integrity, and availability. These translate into specific functionality requirements that include authentication, access control, audit, encryption, secure communication, and non-repudiation. Such requirements are usually satisfied with a suitable security model or architecture, which is commonly expressed in terms of subjects, objects, and actions that subjects can perform on objects. For example, UNIX has a simple security model. Users are subjects. Files are objects. Whether a subject can read, write, or execute an object depends on whether the subject has permission as expressed by the permissions mode specified for the object. However, at lower levels within the system, the security model is expressed with integers, in terms of UID, GID, and the permission mode. Here, the low-level system mechanisms do not (need to) understand the concept of a user and do not (need to) be involved in how a user is authenticated and what UID and GID they are assigned.

[0286] In one embodiment, the peer-to-peer platform protocols may be compatible with widely accepted transport layer security mechanisms for message-based architectures such as Secure Sockets Layer (SSL) and Internet Protocol Security (IPSec). However, secure transport protocols such as SSL and IPSec only provide the integrity and confidentiality of message transfer between two communicating peers. In order to provide secure transfer in multi-hops network, a trust association may be established among all the intermediary peers. Security is compromised if anyone of the communication links is not secured.

[0287] The peer-to-peer platform security model may be implemented to provide a P2P web of trust. The web of trust may be used to exchange public keys among its members. Each peer group policy may permit some members to be trusted to the extent that they have the authority to sign public keys for other members as well as to do things like authenticate, add new members, and remove or revoke membership.

[0288] Embodiments may implement security classes for the RSA public key exchange, the RC4 byte stream cipher, and the SHA-1 hash algorithm, among others. These classes may enable privacy by the means of a P2P TLS implementation; integrity with signed hashes; non-repudiation using the web of trust; and MACs for data authenticity. Combinations of these classes may form security suites, and the peer-to-peer platform provides the mechanism to add new customized suites as required.

[0289] In some embodiments, for peer group authentication a separate Pluggable Authentication Module (PAM) may be provided. Embodiments may provide anonymous or guest login, and login with user name and password. A login session may be in clear or cipher-text as per the peer group security policy.

[0290] The security module may be available to the core level, and thus services, applications and advanced services and applications may plug in their own security components and protocols. For example, the web of trust may be defined by a policy that requires authorized peer group members to be well-known certificate authorities, and that peers exchange X509v3 CA signed certificates.

[0291] Given that the peer-to-peer platform is defined around the concepts of peers and peer groups, one embodiment may include a security architecture in which peer IDs and group IDs are treated as low-level subjects (just like

UID and GID), codats are treated as objects (just like files), and actions are those operations on peers, peer groups, and codats.

[0292] The term "codat" as used herein refers to any computer content - code, data, applications, or other collection of computer representable resources. The peer-to-peer protocol preferably does not distinguish among different types of resources that can be stored on a computer and shared among peers in a peer group. Examples of "codat" include text files, photographs, applets, executable files, serialized Java objects, SOAP messages, etc. Codats are the elementary unit of information that is exchanged among peers. In this embodiment, given that codats may have arbitrary forms and properties, it may not be clear what sets of actions should be defined for them. In one embodiment, the codats may carry or include definitions of how they should be accessed. Such codats are analogous to objects, which define for themselves access methods others can invoke.

[0293] One or more of several other characteristics of the peer-to-peer platform may further affect the security requirements of the peer-to-peer platform. In one embodiment, the peer-to-peer platform may be focused on mechanisms and not policy. For example, UUIDs are used throughout, but they by themselves have no external meaning. Without additional naming and binding services, UUIDs are just numbers that do not correspond to anything like a user or a principal. Therefore, the peer-to-peer platform preferably does not define a high-level security model such as information flow, Bell-LaPadula, or Chinese Wall. In one embodiment, when UUIDs are bound to external names or entities to form security principals, authenticity of the binding may be ensured by placing in the data field security attributes, for example, digital signatures that testify to the trustworthiness of the binding. Once this binding is established, authentication of the principal, access control based on the principal as well as the prevailing security policy, and other functions such as resource usage accounting may be performed.

[0294] The peer-to-peer platform is preferably neutral to cryptographic schemes and security algorithms. As such, the peer-to-peer platform preferably does not mandate any specific security solution. In such cases, a framework may be provided where different security solutions can be plugged in. In one embodiment, hooks and placeholders may be provided so that different security solutions may be implemented. For example, every message may have a designated credential field that may be used to place security-related information. In one embodiment, exactly how to interpret such information is not defined in the peer-to-peer platform, and may be left to services and applications.

[0295] In one embodiment, the peer-to-peer platform may sometimes satisfy security requirements at different levels of the system. To allow maximum flexibility and avoid redundancy, the peer-to-peer platform preferably does not force a particular implementation on developers. Instead, preferably, enhanced platforms based on the peer-to-peer platform may provide the appropriate security solutions to their targeted deployment environment. To illustrate the last point, two security concerns (communications security and anonymity) are examined.

[0296] Peers communicate through pipes. As an example, suppose both confidentiality and integrity in the communications channel are desired. In one embodiment, Virtual Private Networks (VPNs) may be used to move all network traffic. In one embodiment, a secure version of the pipe may be created, similar to a protected tunnel, such that any message transmitted over this pipe is automatically secured. In one embodiment, regular communications mechanisms may be used, and specific data payloads may be protected with encryption techniques and digital signatures. Embodiments of the peer-to-peer platform may accommodate one or more of these and other possible solutions.

[0297] Anonymity does not mean the absence of identity. Indeed, sometimes a certain degree of identification is unavoidable. For example, a cell phone number or a SIM card identification number cannot be kept anonymous, because it is needed by the phone company to authorize and set up calls. As another example, the IP number of a computer cannot be hidden from its nearest gateway or router if the computer wants to send and receive network traffic. In general, anonymity can be built on top of identity, but not vice versa. There may be multiple ways to ensure anonymity. In the examples above, it is difficult to link a prepaid SIM card sold over the retail counter for cash to the actual cell phone user. Likewise, a cooperative gateway or router may help hide the computer's true IP address from the outside world by using message relays or NAT (Network Address Translation).

[0298] In one embodiment, a peer-to-peer platform-based naming service may bind a peer to a human user. The user's anonymity may be ensured through the naming service, or the authentication service, or a proxy service, or any combination of these. The peer-to-peer platform is preferably independent of the solution chosen by a particular application.

[0299] At many places, the peer-to-peer platform is preferably independent of specific security approaches. In one embodiment, the peer-to-peer platform may provide a comprehensive set of security primitives to support the security solutions used by various peer-to-peer platform services and applications. Embodiments of the peer-to-peer platform may provide one or more security primitives including, but not limited to:

- A simple crypto library supporting hash functions (e.g., MD5), symmetric encryption algorithms (e.g., RC4), and asymmetric crypto algorithms (e.g., Diffie-Hellman and RSA).
- An authentication framework that is modeled after PAM (Pluggable Authentication Module, first defined for the UNIX platform and later adopted by the Java security architecture).

- A simple password-based login scheme that, like other authentication modules, can be plugged into the PAM framework.
- A simple access control mechanism based on peer groups, where a member of a group is automatically granted access to all data offered by another member for sharing, whereas non-members cannot access such data.
- 5 • A transport security mechanism that is modeled after SSL/TLS, with the exception that it is impossible to perform a handshake, a crypto strength negotiation, or a two-way authentication on a single pipe, as a pipe is unidirectional.
- The demonstration services called InstantP2P and CMS (content management service) also make use of additional security features provided by the underlying Java platform.

10 **[0300]** In one embodiment, peers, configurations, peer groups, and pipes form the backbone of the peer-to-peer platform. Security in some embodiments of the peer-to-peer platform may use credentials and authenticators (code (e.g. computer-executable instructions) that may be used to receive messages that either request a new credential or request that an existing credential be validated). A credential is a token that when presented in a message body is used to identify a sender and can be used to verify that sender's right to send the message to the specified endpoint and other associated capabilities of the sender. The credential is an opaque token that must be presented each time a message is sent. The sending address placed in the message envelope may be crosschecked with the sender's identity in the credential. In one embodiment, each credential's implementation may be specified as a plug-in configuration, which allows multiple authentication configurations to co-exist on the same network.

15 **[0301]** Preferably, all messages include, at a minimum, a peer group credential that identifies the sender of the message as a full member peer in the peer group in good standing. Membership credentials may be used that define a member's rights, privileges, and role within the peer group. Content access and sharing credentials may also be used that define a member's rights to the content stored within the group.

20 **[0302]** In one embodiment, the peer-to-peer platform may provide different levels of security. In one embodiment, APIs may be provided to access well known security mechanisms such as RCA. In one embodiment, the peer-to-peer platform may provide a distributed security mechanism in a peer-to-peer environment. In one embodiment, this distributed security may not depend on certificates administered by a central authority. The distributed security mechanism may allow a peer group "web of trust" to be generated. In the distributed security mechanism, peers may serve as certificate authorities (security peers). Each peer group may include one or more peers that may serve as a certificate authority in the group. In one embodiment, the creator of a peer group may become the default security authority in the group. In one embodiment, if there is more than one creator, the creator peers may choose one of the peers to be the security authority in the group. In one embodiment, the peer or peers that create a peer group may define the security methods that are to be used within the group (anywhere from no security to high levels of security). In one embodiment, more than one peer in a peer group may serve as a security peer. Since peers are not guaranteed to be up at all times, having multiple security peers in a peer group may help insure that at least one security peer is available at all times. In one embodiment, the peer group's certificate peer may verify keys to provide a weak level of trust. In one embodiment, peer-to-peer platform advertisements may include information to describe the security mechanism (s) to be used in a peer group. For example, the advertisement may include information to do public key exchange, information to indicate what algorithms are to be used, etc. The advertisement may also include information that may be used to enforce secure information exchange on pipes (e.g. encryption information).

30 **[0303]** In one embodiment, peer group security may establish a "social contract". The role of security is distributed across peer groups, and across members of peer groups, that all agree to participate by the rules. A peer group may establish the set of rules by which security in the group is enforced. A peer may join the peer group with a low level of security clearance (low trust). If the peer stays in the group and behaves (follows the rules), the peer may build up its level of trust within the group, and may eventually be moved up in its security level. Within peer groups operating under a social contract, certificates and/or public keys may be exchanged without the participation of a strict certificate authority; i.e. the members may exchange certificates based upon their trust in each other. In one embodiment, a peer group may use an outside challenge (e.g. a secret group password) that may be encrypted/decrypted with public/private keys, as a method to protect and verify messages within the group. In one embodiment, peer groups may be configured to use other types of security, including a high level of security, for example using a strict certificate authority, and even no security. In one embodiment, peer-to-peer platform messages exchanged within a group may have a "placeholder" for security credentials. This placeholder may be used for different types of credentials, depending upon the security implementation of the particular group. In one embodiment, all peer-to-peer messages within the group may be required to have the embedded credential. One embodiment may support private secure pipes.

55 Peer-to-peer platform Firewalls and Security

[0304] The peer-to-peer platform may provide one or more methods for traversing firewalls. Figure 20 illustrates traversing a firewall 248 in a virtual private network when access is initiated from outside only according to one em-

bodiment. Peers 200 on either side of the firewall 248 may each belong to one or more peer groups. In one embodiment, entry may be restricted to peers 200 with access privileges. In this example, peers 200A and 200B have access privileges, but peer 200C does not. Thus, peers 200A and 200B may access peers 200D and 200E through firewall 248. In one embodiment, HTTP "tunnels" may be used, with proxies 246 in the "DMZ" of the firewall 248.

5 **[0305]** Figure 21 illustrates email exchange through a firewall 248 via an email gateway 260 according to one embodiment. In this example, peers 200A and 200B outside the firewall 248 may exchange messages to peers 200C and 200D via the email gateway 260. In one embodiment, there may be an SMTP (Simple Mail Transfer Protocol) service 262 on each peer 200. In one embodiment, 100% peer-to-peer access may not be guaranteed. In one embodiment, inside the firewall 248, mail account administration may impose restrictions. In one embodiment, email addresses may

10 not be required for all peers 200 outside of the firewall 248.
[0306] Figure 22 illustrates several methods of traversing a firewall 248 when access is initiated from the inside according to one embodiment. One or more peers 200 may be inside the firewall 248, and one or more peers 200 may be outside the firewall 248. In one embodiment, each peer 200 that needs to traverse firewall 248 may include a mini-HTTP server. In this embodiment, an HTTP proxy may be used to provide peer-to-peer HTTP tunnels 264 through

15 firewall 248. In one embodiment, Secure Shell (SSH) tunnels 266 may be used to traverse firewall 248. One embodiment may support SOCKS connections 268 if SOCKS is supported in the firewall 248. SOCKS is typically used to telnet/ftp to the "outside". Other embodiments may include other methods of traversing firewalls.

[0307] In one embodiment, peer-to-peer platform core protocols may be used for firewall traversal. In one embodiment, the impact on the peer-to-peer protocol core may be minimized in the traversal method. In one embodiment, 20 peers preferably use the "pure" core protocols for traversal whenever possible. In embodiments where the core protocols need to be extended for traversal, a "divide and conquer" technique is preferably used. In a divide and conquer technique, any new configurations (policies) are preferably isolated behind the firewall. A proxy or proxies may then be used to mediate with and bridge to the core protocols.

[0308] Preferably, peers on either side of the firewall may initiate peer group contact with full peer-to-peer protocol 25 implementation including, but not limited to, the ability to initiate peer group discovery, the ability to join /leave peer groups, and the ability to create end-to-end pipes (cipher text data exchange when required).

[0309] Figure 23 illustrates one embodiment of a peer-to-peer platform proxy service 270, and shows various aspects of the operation of the proxy service. One or more peers 200 may be inside a firewall 248, and one or more peers 200 may be outside the firewall 248. Peer-to-peer platform proxy service 270 is also shown outside the firewall 248. Proxy 30 service 270 may be used to enable peer 200 and peer group contact across firewall 248. Firewall 248 may include an email gateway 260. In one embodiment, the proxy service 270 may be used to bridge peer-to-peer platform protocols 272 with HTTP 274, email 276 and/or SOCKS 278. The proxy service 270 may allow peers 200 to send requests to communicate across firewall 248. Through the proxy service 270, peer-to-peer platform messages may be posted for delivery across the firewall 248. In one embodiment, the proxy service 270 may allow secure pipes to be established

35 across the firewall 248 as necessary.

[0310] Figure 24 illustrates a method of using a proxy service for peer group registration according to one embodiment. The proxy service may permit firewall-independent peer group membership. Three peer regions 212 are shown, with two (region 212A and 212B) on one side of firewall 248 and one (region 212C) on the other side of firewall 248. A peer group 210 may be established that extends across the firewall 248 into regions 212A, 212B and 212C. One or 40 more peers 200 in each region 212 may be members of the peer group 210.

[0311] Figure 25 illustrates peer group registration across a firewall according to one embodiment. Peer region 212A is shown outside of a firewall 248 and peer region 212B is behind the firewall 248. Peer region 212A includes a peer-to-peer platform proxy service 270 and several peers 200. In one embodiment, a peer 200 may be serving as a proxy peer that provides the proxy service 270. Peer region 212B includes several peers 200 behind the firewall 248. At 45 some point, peer 200D in peer region 212B may form a peer group 210. An advertisement for the peer group 210 may be registered on the proxy service 270 in the region 212A. One or more peers 200 in region 212A may be notified of the newly registered peer group 200 by the proxy service 270. In one embodiment, the proxy service may also notify other known peer-to-peer platform proxy services in this or other regions 212, who in turn may notify other proxy services, and so on. Peers 200 in region 212A may then apply for membership in peer group 200.

50 **[0312]** Figure 26 illustrates a method of providing peer group membership through a peer-to-peer platform proxy service according to one embodiment. Peer regions 212A and 212B are shown outside of a firewall 248, and peer region 212C is behind the firewall 248. The two peer group regions 212 outside the firewall 248 each include a proxy service 270. At least one of the peers (peer 200F, in this example) in region 212C behind the firewall belongs to a peer group 210. The peer group 210 may be registered with the proxy services 270 in the regions 212A and 212B outside 55 the firewall 248. A peer 200 in either of the regions outside the firewall may join the peer group 200 by proxy through the proxy service 270 in its region 212. Peers 200 in the regions 212 outside the firewall 248 that are members of the peer group 210 may also leave the peer group 210 through the proxy service 270. Membership information (e.g. included in peer group advertisements) for the peer group 200 may be synchronized on all known proxy services 270

outside the firewall 248. In one embodiment, a proxy service 270 may be a member peer of all locally registered peer groups 200.

[0313] Several levels of authentication may be provided in one or more embodiments of the peer-to-peer platform. Anonymous login may be provided in one embodiment. In one embodiment, a plain text login (user or user and password) may be provided. In one embodiment, login with privacy may be provided. In this embodiment, public key exchange may be used and/or a symmetric master key. The login process preferably returns a credential to the joining peer so that the peer may bypass the login process until the credential expires. One embodiment may provide a public key chain that may be used by registered users to eliminate public key exchanges and thus provides unauthenticated access. One embodiment may provide secure public key exchange with signed certificates.

[0314] Figures 27A and 27B illustrate a method of providing privacy in the peer-to-peer platform according to one embodiment. Figure 27A shows a peer region 212 with peers 200A and 200B and a peer-to-peer platform proxy service 270. Peers 200A and 200B may fetch and cache public keys from a public key chain 280 of the proxy service 270. The cached public keys preferably have expiration dates. Peers 200A and/or 200B may compute a master secret key for one or more of the public keys. Using the keys, cipher text may be exchanged between peers 200A and 200B in privacy as illustrated in Figure 27B.

[0315] The peer-to-peer platform may include one or more methods for providing data integrity in the peer-to-peer environment. These methods may be used to insure that what is sent is what is received. One embodiment may use a standard hash on data (e.g. Secure Hash Algorithm (SHA-1) as defined by the Secure Hash Standard of the Federal Information Processing Standards Publication 180-1). A weak form and/or a strong form may be used in embodiments. In one embodiment, the weak form may use a public key ring and symmetric master to sign data. This method may work best between two peers each having the other's public key. In one embodiment, the strong form may use a symmetric key algorithm such as RSA (Rivest-Shamir-Adleman) and certificate authorities. In one embodiment, the peer-to-peer platform may provide a proxy public certificate authority service. The authority service may create, sign and distribute certificates (e.g. X509 certificates) for all peers on a public key chain. The proxy service's public key is preferably resident on each proxied peer. Other embodiments may utilize other integrity methods.

[0316] Figures 28A and 28B illustrate one embodiment of a method for using a peer-to-peer platform proxy service as a certificate authority. Figure 28A illustrates a peer region 212 with several peers 200 and a proxy service 270. The proxy service 270 may distribute signed certificates in response to peer requests as required. The peers 200 may validate the proxy service 270 signature using a proxy service public key. As illustrated in Figure 28B, when exchanging content with other peers 200, a peer 200 may sign the content with the destination peer's public key and distribute the cipher text.

Bootstrapping mechanism

[0317] In the absence of an application, the peer-to-peer platform preferably provides a mechanism that may be used to discover basic core abstractions (e.g. peer, peer groups, advertisements, pipes). This basic mechanism is needed for bootstrapping a system, and so may be referred to as a bootstrapping mechanism. For example, if a user just downloaded a binary image that enables a device to become a peer in a peer-to-peer network that implements the peer-to-peer platform, the bootstrapping mechanism may be used to discover core abstractions since the "fresh" system may not have knowledge of or access to higher-level services.

[0318] The tasks of searching, discovering, and/or routing in a peer-to-peer network may be complicated. There are many different types of content, and there may not be a generic to best accomplish those tasks for all types of content. Therefore, letting an application or higher-level service perform these high-level search may be preferable, while providing simple, small, mechanisms for bootstrapping peer-to-peer platform-enabled applications.

[0319] The policies and/or protocols used by the core in order to achieve this bootstrapping are preferably as simple as possible and preferably may be implemented and used on a wide variety of platforms (e.g. PDAs, pagers, smart appliances, laptops, workstations, clusters of servers, etc.) and in a variety of network topologies. For example, some peers may not use TCP/IP, and some may not be connected to the Internet. The bootstrapping mechanism may be used as a fallback mechanism when nothing else is useable (e.g. in case of a failure of higher level services). The bootstrapping mechanism is preferably highly configurable. In one embodiment, configuration "wizards" may be used for automatic configuration of the bootstrapping mechanism.

[0320] In one embodiment, other services (e.g. higher-level services and/or optional services) and applications may take over control of the bootstrapping mechanism. In one embodiment, the core protocols may provide an API or APIs to allow the service and/or application to dynamically teach and/or reconfigure the core policies. In one embodiment, a service or application may dynamically overload (i.e. replace) the core policies. For example, this may be done when the design of the application is so dependant on a specific algorithm that it cannot handle the default core policies.

[0321] Providing the bootstrapping mechanism in the peer-to-peer platform may help to allow the peer-to-peer platform to be used straight "out of the box", and/or to be easily configured and installed, for use with a peer-to-peer

platform-enabled application.

Peer Monitoring and Metering

5 [0322] Peer monitoring may include the capability to closely keep track of a (local or remote) peer's status, to control the behavior of a peer, and to respond to actions on the part of a peer. These capabilities may be useful, for example, when a peer network wants to offer premium services with a number of desirable properties such as reliability, scalability, and guaranteed response time. For example, a failure in the peer system is preferably detected as soon as possible so that corrective actions can be taken. It may be preferable to shut down an erratic peer and transfer its responsibilities to another peer.

10 [0323] Peer metering may include the capability to accurately account for a peer's activities, in particular its usage of valuable resources. Such a capability is essential if the network economy is to go beyond flat-rate services. Even for providers offering flat rate services, it is to their advantage to be able to collect data and analyze usage patterns in order to be convinced that a flat rate structure is sustainable and profitable.

15 [0324] In one embodiment, the peer-to-peer platform may provide monitoring and metering through the peer information protocol, where a peer can query another peer for data such as up time and amount of data handled. Security is important in peer monitoring and metering. In one embodiment, a peer may choose to authenticate any command it receives. In one embodiment, a peer may decide to not answer queries from suspect sources.

20 Peer-to-Peer Platform shell application

[0325] One embodiment of the peer-to-peer platform may include a shell application as a development environment built on top of the platform. In one embodiment, the shell application may provide interactive access to the peer-to-peer platform via a simple command line interface. With the shell, shell scripts may be written. The shell may be executed in a networked environment. A user command in the shell may generate a sequence of message exchanges between a set of peers, with some computation occurring on remote peer nodes, and with the answer being returned to the user of the shell. Using the shell, peer-to-peer core building blocks such as peers, peer groups, pipes, and codats may be manipulated. Codats are units of contents that can hold both code and data. For example, a user, through the shell, can publish, search, and execute codats, discover peers or peer groups, create pipes to connect two peers, and send and receive messages.

30 [0326] In one embodiment, an interpreter in the shell may operate in a loop: it accepts a command, interprets the command, executes the command, and then waits for another command. The shell may display a prompt to notify users that it is ready to accept a new command.

35 [0327] In one embodiment with a Java-based implementation of the peer-to-peer platform, one or more of the shell commands may not be built in per se. The commands may be Java language programs and are dynamically loaded and started by the shell framework when the corresponding commands are typed in. Therefore, adding a new shell command may be performed by writing a program in the Java language.

40 [0328] In one embodiment, the shell may provide a "pipe" capability to redirect a command output pipe into another command input pipe. In one embodiment, shell commands may be given a standard input, output and error pipes that a user can connect, disconnect and reconnect to other shell commands. Commands can support other pipes if needed. In one embodiment of the shell, a user may dynamically disconnect and reconnect pipes between commands, as in the following example:

45 xxxx> cat >p1 myfile
 xxxx> grep <p1 abcd
 xxxx> grep <p1 efgh

[0329] In the above example, the first command "cat >p1 myfile" cats myfile into the output pipe p1. The second command then connects pipe p1 to grep's input pipe and searches for the string abcd. The third command then disconnects p1, redirects it to the new grep command's input pipe and searches for the string efgh.

50 [0330] In one embodiment, the peer-to-peer platform shell supports piping in both directions. A special operator such as "<>" may be used for creating crossing pipes between two commands. For example, with the following command "cmd1 <> cmd2", the output pipe of the first command is connected to the standard input pipe of the second command, and at the same time the output pipe of the second command is connected to the standard input pipe of the first command. Of course, this operator has to be used carefully to avoid infinite data loops.

55 [0331] In one embodiment, applications other than peer-to-peer platform applications may be run from the shell. For content management, MIME type information included with a codat may be used to let local applications associated with well-known content types handle them automatically. The peer-to-peer platform may support the development of

adaptors to allow the execution of external programs with appropriate security safeguards. An adapter may essentially map data and connect applications for remote launches. Some examples might be:

- UNIX® stdio to peer-to-peer platform stdio adapter — such an adaptor may enable piping of peer-to-peer platform commands to UNIX® commands on UNIX® platforms.
- Peer-to-peer platform stream to a media player adapter — such an adaptor may be platform- and application-specific, but may handle any necessary real-time data conversion between a peer-to-peer platform pipe and the format required by the player. These need not be unidirectional. For example, one might adapt the output of a video capture application to become a peer-to-peer platform stream.
- HTML to peer-to-peer platform stdio — such an adaptor may be used to post and get information to/from Web pages to allow peer-to-peer platform peers to interact with existing Web sites. For example, a peer-to-peer platform command can launch a search for titles and prices on Amazon and pipe the results to other peer-to-peer platform services.

Conclusion

[0332] Various aspects of the present invention are summarized in the following numbered clauses 1 - 116:

1. A peer computing system comprising:

a plurality of peer nodes operable to couple to a network;

wherein the plurality of peer nodes are configured to implement a peer-to-peer environment on the network according to a peer-to-peer platform comprising:

a core layer comprising one or more peer-to-peer platform protocols for enabling the plurality of peer nodes to discover each other, communicate with each other, and cooperate with each other to form peer groups and share content in the peer-to-peer environment;

a service layer comprising one or more core services each provided by one or more of the plurality of peer nodes in the peer-to-peer environment, wherein at least a subset of the core services are operable to be used by the plurality of peer nodes in forming and participating in the peer groups, and wherein each of the one or more core services are configured to be accessed by the plurality of peer nodes in accordance with at least one of the one or more peer-to-peer platform protocols; and

an application layer comprising one or more applications each provided by one or more of the plurality of peer nodes in the peer-to-peer environment, wherein each of the one or more applications are configured to be accessed in accordance with at least one of the one or more peer-to-peer platform protocols, and wherein at least a subset of the one or more applications are each configured to access at least one of the one or more core services to perform application tasks in the peer-to-peer environment in accordance with at least one of the one or more peer-to-peer platform protocols.

2. The peer computing system as recited in clause 1, wherein the service layer further comprises one or more other services that are not core services in the peer-to-peer environment.

3. The peer computing system as recited in clause 1, wherein each of the one or more peer-to-peer platform protocols defines one or more advertisement formats for describing and publishing advertisements for resources in the peer-to-peer environment.

4. The peer computing system as recited in clause 3, wherein the resources include one or more of the peer nodes, the peer groups, the content, the core services, other services in the service layer, the applications, pipes, and pipe endpoints, wherein the pipes are communications channels between one or more of the peer nodes, the core services, the other services and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels.

5. The peer computing system as recited in clause 1, wherein at least a subset of the one or more peer-to-peer platform protocols defines one or more message formats configured for use in exchanging messages between the peer nodes in accordance with the particular protocol.

6. The peer computing system as recited in clause 1, wherein the one or more peer-to-peer platform protocols includes one or more of:

a peer discovery protocol for discovering resources in the peer-to-peer environment;

a peer membership protocol for use by the peer nodes in applying for membership in the peer groups;
 a peer resolver protocol for use in sending search queries from one peer group member to another peer group member;

a peer information protocol for enabling the peer nodes to obtain information about capabilities and status of other peer nodes in the peer-to-peer environment;

a pipe binding protocol for use in finding the physical location of pipe endpoints and binding the pipe endpoints, wherein pipes are communications channels between one or more of the peer nodes, the core services and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels;

an endpoint routing protocol for enabling the peer nodes to request peer routing information to reach the other peer nodes; and

a peer rendezvous protocol for enabling peer nodes to propagate query messages to a next set of peer nodes.

7. The peer computing system as recited in clause 6, wherein the resources include one or more of the peer nodes, the peer groups, the content, the core services, other services in the service layer, the applications, pipes, and pipe endpoints, wherein the pipes are communications channels between one or more of the peer nodes, the core services, the other services and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels.

8. The peer computing system as recited in clause 1, wherein the one or more peer-to-peer platform protocols includes a discovery protocol for discovering the peer nodes in the peer-to-peer environment.

9. The peer computing system as recited in clause 8, wherein the one or more peer-to-peer platform protocols define a peer advertisement format configured for use in advertising the peer nodes in the peer-to-peer environment, wherein said discovering the peer nodes returns one or more peer advertisements for the discovered peer nodes formatted in accordance with the peer advertisement format.

10. The peer computing system as recited in clause 1, wherein the one or more peer-to-peer platform protocols includes a discovery protocol for discovering the peer groups in the peer-to-peer environment.

11. The peer computing system as recited in clause 10, wherein the one or more peer-to-peer platform protocols define a peer group advertisement format configured for use in advertising the peer groups in the peer-to-peer environment, wherein said discovering the peer groups returns one or more peer group advertisements formatted in accordance with the peer group advertisement format.

12. The peer computing system as recited in clause 1, wherein the one or more peer-to-peer platform protocols includes a discovery protocol for enabling the peer nodes to discover and exchange content in the peer-to-peer environment.

13. The peer computing system as recited in clause 12, wherein the one or more peer-to-peer platform protocols define a content advertisement format configured for use in advertising the content in the peer-to-peer environment, wherein said discovering content returns one or more content advertisements formatted in accordance with the content advertisement format.

14. The peer computing system as recited in clause 1, wherein the one or more peer-to-peer platform protocols include a discovery protocol for discovering pipes in the peer-to-peer environment, wherein the pipes are communications channels between one or more of the peer nodes, the core services and the applications in the peer-to-peer environment.

15. The peer computing system as recited in clause 14, wherein the one or more peer-to-peer platform protocols define a pipe advertisement format configured for use in advertising pipes in the peer-to-peer environment, wherein said discovering pipes returns one or more pipe advertisements formatted in accordance with the pipe advertisement format.

16. The peer computing system as recited in clause 1, wherein the one or more peer-to-peer platform protocols include a discovery protocol for discovering pipe endpoints in the peer-to-peer environment, wherein the pipes are communications channels between one or more of the peer nodes, the core services and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels.

17. The peer computing system as recited in clause 16, wherein the one or more peer-to-peer platform protocols define an endpoint advertisement format configured for use in advertising endpoints in the peer-to-peer environment, wherein said discovering endpoints returns one or more endpoint advertisements formatted in accordance with the endpoint advertisement format.

18. The peer computing system as recited in clause 1, wherein the one or more peer-to-peer platform protocols includes a discovery protocol for discovering the core services and other services provided by the peer nodes in the peer-to-peer environment.

19. The peer computing system as recited in clause 18, wherein the one or more peer-to-peer platform protocols define a service advertisement format configured for use in advertising the core services and the other services provided by the peer nodes in the peer-to-peer environment, wherein said discovering the core services and the other services returns one or more service advertisements formatted in accordance with the service advertisement format.

20. The peer computing system as recited in clause 1, wherein the one or more peer-to-peer platform protocols includes a peer membership protocol for use by the peer nodes in applying for membership in one or more of the peer groups.

21. The peer computing system as recited in clause 1, wherein the one or more peer-to-peer platform protocols include a peer resolver protocol for use in sending generic search queries from one peer node to one or more other peer nodes in the peer-to-peer environment.

22. The peer computing system as recited in clause 21, wherein the search queries are sent to one or more services configured to perform searches as specified by the search queries and to generate responses to the search queries, wherein the one or more services are each hosted by one of the one or more other peer nodes.

23. The peer computing system as recited in clause 22, wherein each of the one or more services is configured to find one or more of peer, peer group, content, service, application, pipe, and pipe endpoint information in accordance with each particular search query received by the particular service handler, wherein the pipes are communications channels between one or more of the peer nodes, the core services, other services in the service layer, and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels.

24. The peer computing system as recited in clause 1, wherein the one or more peer-to-peer platform protocols include a pipe binding protocol for use in finding the physical location of a pipe endpoint and in binding to the pipe endpoint, wherein pipes are communications channels between one or more of the peer nodes, the core services, other services in the service layer, and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels.

25. The peer computing system as recited in clause 1, wherein the one or more peer-to-peer platform protocols include an endpoint routing protocol for enabling the peer nodes to request peer routing information to reach other peer nodes.

26. The peer computing system as recited in clause 25, wherein, in said requesting peer routing information, the peer nodes are configured to use the endpoint routing protocol to send route query request messages formatted in accordance with the endpoint routing protocol to one or more router peers to request the peer routing information.

27. The peer computing system as recited in clause 26, wherein each of the router peers is configured to cache route information for one or more routes in the peer-to-peer environment, and wherein each of the router peers is further configured to return route information for a particular route specified by a particular route query request message if the route information for the particular route is cached by the particular router peer.

28. The peer computing system as recited in clause 27, wherein each of the router peers is further configured to forward the route query request message to other router peers if the route information for the particular route is not cached by the particular router peer.

29. The peer computing system as recited in clause 1, wherein the one or more peer-to-peer platform protocols includes a peer information protocol for enabling the peer nodes to obtain information about capabilities and status of other peer nodes in the peer-to-peer environment.

30. The peer computing system as recited in clause 1, wherein each peer group is a collection of cooperating member peer nodes that provides a common set of services to the member peer nodes in the peer-to-peer environment.

31. The peer computing system as recited in clause 30, wherein the common set of services on at least a subset of the peer groups includes one or more of a discovery service, a membership service, an access service, a pipe service, a resolver service and a monitoring service, wherein pipes are communications channels between one or more of the peer nodes, the core services, other services in the service layer, and the applications in the peer-to-peer environment.

32. The peer computing system as recited in clause 30, wherein the peer-to-peer platform protocols include a discovery protocol, wherein the common set of services on at least a subset of the peer groups includes a discovery service for use by member peer nodes in said peer group to discover advertised resources including peer nodes and peer groups in the peer computing system in accordance with the discovery protocol.

33. The peer computing system as recited in clause 30, wherein the peer-to-peer platform protocols include a membership protocol, wherein the common set of services on at least a subset of the peer groups includes a membership service for use by member peer nodes in said peer group to reject or accept group membership applications in accordance with the membership protocol.

34. The peer computing system as recited in clause 30, wherein the common set of services includes one or more user-defined services.

35. The peer computing system as recited in clause 1, wherein each of the plurality of peer nodes includes a unique identifier configured for use in distinguishing each peer node from the other peer nodes in the peer-to-peer environment.

36. A peer node comprising:

one or more network interfaces for coupling to a network;
a memory comprising program instructions, wherein the program instructions are executable within the peer node to implement, according to a peer-to-peer platform:

a core layer comprising one or more peer-to-peer platform protocols for enabling the peer node to discover other peer nodes, communicate with the other peer nodes, and cooperate with the other peer nodes to form peer groups and share content in a peer-to-peer environment on the network;

a service layer comprising one or more core services in the peer-to-peer environment, wherein at least a subset of the core services are operable to be used by the peer node and the other peer nodes in forming and participating in the peer groups, and wherein each of the one or more core services are configured to be accessed in accordance with at least one of the one or more peer-to-peer platform protocols; and

an application layer comprising one or more applications, wherein each of the one or more applications are configured to be accessed by the peer node and the other peer nodes in accordance with at least one of the one or more peer-to-peer platform protocols, and

wherein at least a subset of the one or more applications are each configured to access at least one of the one or more core services to perform application tasks in the peer-to-peer environment in accordance with at least one of the one or more peer-to-peer platform protocols.

37. The peer node as recited in clause 36, wherein the service layer further comprises one or more other services that are not core services in the peer-to-peer environment.

38. The peer node as recited in clause 36, wherein the program instructions are further executable to host one or more services in a peer group in which the peer node is a member peer, wherein other member peer nodes access the hosted services from the peer node.

39. The peer node as recited in clause 36, wherein the program instructions are further executable to publish advertisements for resources in the peer-to-peer environment using one or more advertisement formats defined by the peer-to-peer platform protocols, wherein the resources include one or more of the peer nodes, the peer groups, content, the core services, other services in the services layer, the applications, pipes and pipe endpoints, wherein the pipes are communications channels between one or more of the peer nodes, the core services, the other services and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels.

40. The peer node as recited in clause 36, wherein the program instructions are further executable to send messages to and receive messages from the other peer nodes in the peer-to-peer environment using one or more message formats each defined by one of the one or more peer-to-peer platform protocols.

41. The peer node as recited in clause 36, wherein the one or more peer-to-peer platform protocols includes one or more of:

a peer discovery protocol for use by the peer node in discovering resources in the peer-to-peer environment, wherein the resources include one or more of the peer nodes, the peer groups, content, services, pipes and pipe endpoints;

a peer membership protocol for use by the peer node in applying for membership in the peer groups;

a peer resolver protocol for use in sending search queries from the peer node to other peer nodes in the peer-to-peer environment;

a peer information protocol for enabling the peer node to obtain information about capabilities and status of the other peer nodes;

a pipe binding protocol for use by the peer node in finding the physical location of pipe endpoints and binding the pipe endpoints, wherein pipes are communications channels between one or more of the peer nodes, the core services and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels;

an endpoint routing protocol for enabling the peer node to request peer routing information to reach one or more of the other peer nodes in the peer-to-peer environment; and

a peer rendezvous protocol for enabling peer nodes to propagate query messages to a next set of peer nodes.

42. The peer node as recited in clause 36, wherein the one or more peer-to-peer platform protocols includes a discovery protocol, wherein the program instructions are further executable to discover resources in the peer-to-peer environment in accordance with the discovery protocol, wherein, in said discovering the resources, the program instructions are further executable to receive one or more advertisements for the discovered resources formatted in accordance with the discovery protocol.

43. The peer node as recited in clause 42, wherein the resources include one or more of the peer nodes, the peer groups, the content, the core services, other services in the service layer, the applications, pipes, and pipe endpoints, wherein the pipes are communications channels between one or more of the peer nodes, the core services, the other services and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels.

44. The peer node as recited in clause 36, wherein the one or more peer-to-peer platform protocols includes a peer membership protocol, wherein the program instructions are further executable to apply for membership in one or more of the peer groups in accordance with the peer membership protocol.

45. The peer node as recited in clause 36, wherein the one or more peer-to-peer platform protocols includes a peer resolver protocol, wherein the program instructions are further executable to send generic search queries to one or more of the other peer nodes in accordance with the peer resolver protocol.

46. The peer node as recited in clause 36, wherein the one or more peer-to-peer platform protocols include a pipe binding protocol, wherein pipes are communications channels between one or more of the peer nodes, the core services, other services in the service layer, and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels, and wherein the program instructions are further executable to:

find the physical location of a pipe endpoint in accordance with the pipe binding protocol; and
bind to the pipe endpoint in accordance with the pipe binding protocol.

47. The peer node as recited in clause 36, wherein the one or more peer-to-peer platform protocols include an endpoint routing protocol, wherein the program instructions are further executable to request peer routing information to the other peer nodes in the peer-to-peer environment in accordance with the endpoint routing protocol.

48. The peer node as recited in clause 36, wherein the one or more peer-to-peer platform protocols includes a peer information protocol, wherein the program instructions are further executable to obtain information about capabilities and status of the other peer nodes in the peer-to-peer environment in accordance with the peer information protocol.

49. The peer node as recited in clause 36, wherein the peer node is a member peer node in a peer group, wherein the peer group is a collection of cooperating member peer nodes that provides a common set of services to the member peer nodes.

50. The peer node as recited in clause 49, wherein the peer-to-peer platform protocols include a discovery protocol, wherein the common set of services provided by the peer group includes a discovery service, wherein the program instructions are further executable to discover advertised resources including the other peer nodes and the peer groups in the peer-to-peer environment using the discovery service in accordance with the discovery protocol.

51. The peer node as recited in clause 49, wherein the peer-to-peer platform protocols include a membership protocol, wherein the common set of services includes a membership service, wherein the program instructions are further executable to reject or accept group membership applications using the membership service in accordance with the membership protocol.

52. The peer node as recited in clause 36, wherein the peer node includes a unique identifier configured to distinguish the peer node from the other peer nodes in the peer-to-peer environment

53. A peer node comprising:

one or more network interfaces for coupling to a network;
a memory comprising program instructions, wherein the program instructions are executable within the peer node to discover and access an instance of a service on one of a plurality of peer nodes,

wherein the one of the plurality of peer nodes is local to a network location of the peer node on the network, wherein the plurality of peer nodes each host an instance of the same service, and wherein said discovering and accessing the instance of the service are performed in accordance with one or more peer-to-peer platform protocols;

wherein the peer node is configured to move from the network location to a different network location;

wherein the program instructions are further executable within the peer node to discover and access a different instance of the service on a different one of the plurality of peer nodes, wherein the different one of the plurality of peer nodes is local to the different network location, and wherein said discovering and accessing the different instance of the service are performed in accordance with the one or more peer-to-peer platform protocols.

54. The peer node as recited in clause 53, wherein the peer node includes a unique identifier of the peer node, wherein the unique identifier distinguishes the peer node from the other peer nodes on the network. wherein the program instructions are further executable to provide the unique identifier to the different instance of the service, and wherein the different instance of the service is operable to recognize the peer node using the unique identifier and to route information to the peer node at the different network location.

55. A peer computing system comprising:

a plurality of peer nodes, wherein the plurality of peer nodes each implement one or more peer-to-peer platform protocols for enabling the plurality of peer nodes to host and access services in a peer-to-peer environment; at least a subset of the plurality of peer nodes that each host an instance of a service;

wherein each of the at least a subset of the plurality of peer nodes is operable to provide access to an instance of the service hosted by the particular peer node to a different one of the plurality of peer nodes at a network location, wherein the particular peer node is local to the network location;

wherein the different one of the plurality of peer nodes is operable to:

move to a different network location; and
provide a unique identifier to the instance of the service hosted by the particular peer node,

wherein the unique identifier distinguishes the different one of the plurality of peer nodes from the other peer nodes on the network;

wherein the instance of the service is operable to recognize the different one of the plurality of peer nodes using the unique identifier and to route information provided by the service to the different one of the plurality of peer nodes at the different network location.

56. A peer node comprising:

one or more network interfaces for coupling to a network;
a memory comprising program instructions, wherein the program instructions are executable within the peer node to discover and access an instance of a service on one of one or more peer nodes,

wherein the one of the one or more peer nodes is local to a network location of the peer node on the network, wherein the one or more peer nodes each host an instance of the same service, and wherein said discovering and accessing the instance of the service are performed in accordance with one or more peer-to-peer platform protocols;

wherein the peer node is configured to move from the network location to a different network location;
wherein the program instructions are further executable within the peer node to:

discover and access the same instance of the service on the one of the one or more peer nodes,

wherein said discovering and accessing the same instance of the service are performed in accordance with the one or more peer-to-peer platform protocols;

provide a unique identifier for the peer node to the instance of the service. wherein the unique identifier distinguishes the peer node from the other peer nodes on the network, and wherein the instance of the service is operable to recognize the peer node using the unique identifier and to route information provided by the service to the peer node at the different network location.

57. A peer computing system comprising:

a plurality of peer nodes, wherein the plurality of peer nodes each implement one or more peer-to-peer platform protocols for enabling the plurality of peer nodes to discover and access contents in a peer-to-peer environment;

at least a subset of the plurality of peer nodes that each include an instance of a content;

wherein each of the plurality of peer nodes is configured to:

discover and access an instance of the content on one of the at least a subset of the plurality of peer nodes, wherein the one of the at least a subset of the plurality of peer nodes is local to a network location of the particular peer node on the network, wherein said discovering and accessing the instance of the content is performed in accordance with the one or more peer-to-peer platform protocols;
 move from the network location to a different network location;
 discover and access a different instance of the content on a different one of the at least a subset of the plurality of peer nodes, wherein the one of the at least a subset of the plurality of peer nodes is local to the different network location, wherein said discovering and accessing the different instance of the content are performed in accordance with the one or more peer-to-peer platform protocols.

58. A peer node comprising:

one or more network interfaces for coupling to a network;
 a memory comprising program instructions, wherein the program instructions are executable within the peer node to discover and access an instance of a content on one of a plurality of peer nodes,

wherein the one of the plurality of peer nodes is local to a network location of the peer node on the network, wherein the plurality of peer nodes each host an instance of the same content, and wherein said discovering and accessing the instance of the service are performed in accordance with one or more peer-to-peer platform protocols;

wherein the peer node is configured to move from the network location to a different network location;

wherein the program instructions are further executable within the peer node to discover and access a different instance of the content on a different one of the plurality of peer nodes, wherein the different one of the plurality of peer nodes is local to the different network location, and wherein said discovering and accessing the different instance of the content are performed in accordance with the one or more peer-to-peer platform protocols.

59. A peer computing system comprising:

a plurality of peer nodes operable to couple to a network;
 means for the peer nodes to discover each other, communicate with each other, and cooperate with each other to form peer groups and share content in a peer-to-peer environment on the network;
 means for the peer nodes to provide, discover and access one or more services in the peer-to-peer environment, wherein at least a subset of the services are core services operable to be used by the plurality of peer nodes in forming and participating in the peer groups; and
 means for the peer nodes to provide, discover and access one or more applications in the peer-to-peer environment; and
 means for at least a subset of the one or more applications to discover and access at least one of the one or more services to perform application tasks in the peer-to-peer environment.

60. The peer computing system as recited in clause 59, further comprising means for the one or more services to discover and access each other in the peer-to-peer environment.

61. The peer computing system as recited in clause 59, further comprising means for describing and publishing resources in the peer-to-peer environment, wherein the resources include one or more of the peer nodes, the peer groups, the content, the services, the applications, pipes, and pipe endpoints, wherein the pipes are communications channels between one or more of the peer nodes, the services and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels.

62. The peer computing system as recited in clause 59, further comprising means for providing communications channels for the peer nodes, the services and the applications to exchange information in the peer-to-peer environment.

63. The peer computing system as recited in clause 59, further comprising means for exchanging messages between the peer nodes in the peer-to-peer environment.

64. The peer computing system as recited in clause 59, further comprising means for discovering resources in the peer-to-peer environment, wherein the resources include one or more of the peer nodes, the peer groups, the content, the services, the applications, pipes and pipe endpoints, wherein the pipes are communications channels between one or more of the peer nodes, the services and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels.

65. The peer computing system as recited in clause 59, further comprising means for the peer nodes to apply for membership in one or more of the peer groups.

66. The peer computing system as recited in clause 59, further comprising means for sending generic search queries from one of the peer nodes to one or more other of the peer nodes.

67. The peer computing system as recited in clause 59, further comprising:

means for finding communications channels between one or more of the peer nodes, the services and the applications in the peer-to-peer environment; and
means for binding to the communications channels.

68. The peer computing system as recited in clause 59, further comprising means for the peer nodes to request peer routing information to reach other peer nodes in the peer-to-peer environment.

69. The peer computing system as recited in clause 59, further comprising means for the peer nodes to obtain information about capabilities and status of other peer nodes in the peer-to-peer environment.

70. The peer computing system as recited in clause 59, wherein the peer groups are collection of cooperating member peer nodes, further comprising means for the peer groups to each provide a common set of services to its member peer nodes.

71. The peer computing system as recited in clause 59, further comprising means for member peer nodes in a peer group to receive and reject or accept group membership applications.

72. The peer computing system as recited in clause 59, further comprising means for distinguishing each peer node from the other peer nodes on the network.

73. A peer computing system comprising:

a plurality of peer nodes configured to couple to a network;
means for the peer nodes to discover each other, communicate with each other, and cooperate with each other to form peer groups and host services in a peer-to-peer environment on the network;

wherein at least a subset of the plurality of peer nodes each hosts an instance of a particular service;

means for each of the plurality of peer nodes to discover and access an instance of a service provided by one of the at least a subset of the plurality of peer nodes, wherein the one of the at least a subset of the plurality of peer nodes is local to a network location of the particular one of the plurality of peer nodes;

wherein each of the plurality of peer nodes is operable to move to a different network location; and

means for each of the plurality of peer nodes to discover and access a different instance of the service provided by a different one of the at least a subset of the plurality of peer nodes, wherein the one of the at least a subset of the plurality of peer nodes is local to the different network location of the particular one of the plurality of peer nodes.

74. The peer computing system of clause 73, further comprising means for the different instance of the service to recognize the particular one of the plurality of peer nodes and to route information provided by the service to the particular one of the plurality of peer nodes at the different network location.

75. A peer computing system comprising:

a plurality of peer nodes configured to couple to a network;
means for the peer nodes to discover each other, communicate with each other, and cooperate with each other to form peer groups and host services in a peer-to-peer environment on the network;

wherein at least a subset of the plurality of peer nodes each hosts an instance of a particular service;

means for each of the plurality of peer nodes to discover and access an instance of a service provided by one of the at least a subset of the plurality of peer nodes, wherein the one of the at least a subset of the plurality of peer nodes is local to a network location of the particular one of the plurality of peer nodes;

wherein each of the plurality of peer nodes is operable to move to a different network location;

means for each of the plurality of peer nodes to access the instance of the service provided by the one of the

at least a subset of the plurality of peer nodes from the different network location of the particular one of the plurality of peer nodes; and
 means for the instance of the service to recognize the particular one of the plurality of peer nodes and to route information provided by the service to the particular one of the plurality of peer nodes at the different network location.

76. A peer computing system comprising:

a plurality of peer nodes operable to couple to a network;
 means for the peer nodes to discover each other, communicate with each other, and cooperate with each other to form peer groups and to share content;

wherein at least a subset of the plurality of peer nodes each hosts an instance of a particular content;

means for each of the plurality of peer nodes to discover and access an instance of a content provided by one of the at least a subset of the plurality of peer nodes, wherein the one of the at least a subset of the plurality of peer nodes is local to a network location of the particular one of the plurality of peer nodes;

wherein each of the plurality of peer nodes is operable to move to a different network location; and

means for each of the plurality of peer nodes to discover and access a different instance of the content provided by a different one of the at least a subset of the plurality of peer nodes, wherein the different one of the at least a subset of the plurality of peer nodes is local to the different network location of the particular one of the plurality of peer nodes.

77. A method for implementing a peer-to-peer environment on a network, the method comprising:

a plurality of peer nodes coupled to a network each implementing a core layer of a peer-to-peer platform, wherein the core layer comprises one or more peer-to-peer platform protocols for enabling the plurality of peer nodes to discover each other, communicate with each other, and cooperate with each other to form peer groups and share content in the peer-to-peer environment;

the plurality of peer nodes each implementing a service layer comprising one or more core services each provided by one or more of the plurality of peer nodes in the peer-to-peer environment, wherein each of the one or more core services are configured to be accessed by peer nodes in the peer-to-peer environment in accordance with at least a subset of the one or more peer-to-peer platform protocols;

the plurality of peer nodes each implementing an application layer comprising one or more applications each provided by one or more of the plurality of peer nodes in the peer-to-peer environment, wherein each of the one or more applications are configured to be accessed in accordance with at least one of the one or more peer-to-peer platform protocols, and wherein at least a subset of the one or more applications are each configured to access at least one of the one or more core services to perform application tasks in the peer-to-peer environment in accordance with at least one of the one or more peer-to-peer platform protocols; and
 at least a subset of the plurality of peer nodes accessing at least a subset of the core services in accordance with at least one of the one or more peer-to-peer platform protocols to form one or more peer groups in the peer-to-peer environment.

79. The method as recited in clause 77, wherein the one or more peer-to-peer platform protocols include a peer membership protocol for joining or forming a peer group with other peer nodes, wherein the one or more core services include a membership service for use by the peer nodes in forming the peer groups and joining the peer groups, wherein the membership service is configured to be accessed by the peer nodes in the peer-to-peer environment in accordance with the membership protocol, the method further comprising one or more of the plurality of peer nodes forming a peer group in the peer-to-peer environment using the membership service.

79. The method as recited in clause 78, further comprising:

another peer node applying for membership in the peer group using the membership service;
 one or more member peer nodes of the peer group determining if the other peer node is qualified for membership in the peer group in response to the application for membership using the membership service; and
 if the member peer nodes determine that the other peer node is qualified for membership in the peer group, the other peer node becoming a member peer node in the peer group.

80. The method as recited in clause 83, wherein the one or more peer-to-peer platform protocols include a discovery protocol for discovering resources in the peer-to-peer environment, and wherein the one or more core services include a discovery service for use by the peer nodes to discover advertised resources in the peer-to-peer environment, wherein the discovery service is configured to be accessed by the peer nodes in the peer-to-peer environment in accordance with the discovery protocol.

81. The method as recited in clause 80, wherein the advertised resources include one or more of the peer nodes, the peer groups, the content, the core services, other services in the service layer, the applications, pipes, and pipe endpoints, wherein the pipes are communications channels between one or more of the peer nodes, the core services, the other services and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels.

82. The method as recited in clause 80, wherein the resources include the peer nodes, the method further comprising:

one of the plurality of peer nodes broadcasting a peer discovery message in the peer-to-peer environment using the discovery service; and
the one of the plurality of peer nodes receiving one or more response messages in response to the peer discovery message, wherein the response messages each include information about a particular peer node, wherein the information is configured for use by the one of the plurality of peer nodes in establishing a connection to the particular peer node; and

wherein the peer discovery message and the one or more response messages are in a format defined by the discovery protocol, and wherein said broadcasting a peer discovery message and said receiving one or more response messages are performed in accordance with the discovery protocol.

83. The method as recited in clause 80, wherein the resources include the peer groups, the method further comprising:

one of the plurality of peer nodes broadcasting a peer group discovery message in the peer-to-peer environment using the discovery service; and
the one of the plurality of peer nodes receiving a peer group response message in response to the peer group discovery message from each of one or more of the peer groups in the peer-to-peer environment, wherein the peer group response messages each include information about a particular peer group, wherein the information is configured for use by the one of the plurality of peer nodes in joining the particular peer group; and

wherein the peer group discovery message and the peer group response message are in a format defined by the discovery protocol, and wherein said broadcasting a peer group discovery message and said receiving a peer group response message are performed in accordance with the discovery protocol.

84. The method as recited in clause 77, further comprising publishing advertisements for resources in the peer-to-peer environment using one or more advertisement formats each defined by one of the one or more peer-to-peer platform protocols.

85. The method as recited in clause 84, wherein the resources include one or more of the peer nodes, the peer groups, the content, the core services, other services in the service layer, the applications, pipes, and pipe endpoints, wherein the pipes are communications channels between one or more of the peer nodes, the core services, the other services and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels.

86. The method as recited in clause 77, further comprising two or more of the plurality of peer nodes exchanging messages in the peer-to-peer environment using one or more message formats each defined by one of the one or more peer-to-peer platform protocols.

87. The method as recited in clause 77, wherein the one or more peer-to-peer platform protocols include a discovery protocol, the method further comprising a peer node discovering resources in the peer-to-peer environment in accordance with the discovery protocol, wherein said discovering the resources comprises the peer node receiving one or more advertisements for the discovered resources formatted in accordance with the peer discovery protocol.

88. The method as recited in clause 87, wherein the resources include one or more of the peer nodes, the peer groups, the content, the core services, other services in the service layer, the applications, pipes, and pipe endpoints, wherein the pipes are communications channels between one or more of the peer nodes, the core services, the other services and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels.

89. The method as recited in clause 77, wherein the one or more peer-to-peer platform protocols includes a peer

membership protocol, the method further comprising one of the plurality of peer nodes applying for membership in one or more of the peer groups in accordance with the peer membership protocol.

90. The method as recited in clause 77, wherein the one or more peer-to-peer platform protocols includes a peer resolver protocol, the method further comprising one of the plurality of peer nodes sending one or more generic search queries to one or more other peer nodes in the peer-to-peer environment in accordance with the peer resolver protocol.

91. The method as recited in clause 77, wherein the one or more peer-to-peer platform protocols include a pipe binding protocol, the method further comprising:

- one of the plurality of peer nodes finding the physical location of a pipe endpoint in accordance with the pipe binding protocol; and
- the peer node binding to the pipe endpoint in accordance with the pipe binding protocol;

wherein pipes are communications channels between one or more of the peer nodes, the core services, other services in the service layer, and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels.

92. The method as recited in clause 77, wherein the one or more peer-to-peer platform protocols include an endpoint routing protocol, the method further comprising one of the plurality of peer nodes requesting peer routing information to other peer nodes in the peer-to-peer environment in accordance with the endpoint routing protocol.

93. The method as recited in clause 77, wherein the one or more peer-to-peer platform protocols include a peer information protocol, the method further comprising one of the plurality of peer nodes obtaining information about capabilities and status of one or more other peer nodes in the peer-to-peer environment in accordance with the peer information protocol.

94. The method as recited in clause 77, wherein each peer group is a collection of cooperating member peer nodes, further comprising each peer group providing a common set of services to the member peer nodes in the peer-to-peer environment.

95. The method as recited in clause 94, wherein the one or more peer-to-peer platform protocols include a discovery protocol, wherein the common set of services includes a discovery service, wherein the discovery service is accessible in accordance with the discovery protocol, the method further comprising one of the member peer nodes in one of the peer groups discovering advertised resources in the peer-to-peer environment using the discovery service.

96. The method as recited in clause 94, wherein the one or more peer-to-peer platform protocols include a membership protocol, wherein the common set of services includes a membership service, wherein the membership service is accessible in accordance with the membership protocol, the method further comprising:

- a peer node not in one of the peer groups applying for membership in the peer group; and
- the member peer nodes of the peer group rejecting or accepting the peer node's group membership application using the membership service.

97. A method comprising:

- a peer node discovering an instance of a service on one of a plurality of peer nodes, wherein the one of the plurality of peer nodes is local to a network location of the peer node on a network, wherein the plurality of peer nodes each host an instance of the same service;
- the peer node accessing the instance of the service;

wherein said discovering and said accessing the instance of the service are performed in accordance with one or more peer-to-peer platform protocols;

- the peer node moving from the network location to a different network location;
- the peer node discovering a different instance of the service on a different one of the plurality of peer nodes, wherein the different one of the plurality of peer nodes is local to the different network location;
- the peer node accessing the different instance of the service; and

wherein said discovering and accessing the different instance of the service are performed in accordance with the one or more peer-to-peer platform protocols.

98. The method as recited in clause 97, further comprising:

the peer node providing a unique identifier for the peer node to the different instance of the service, wherein the unique identifier distinguishes the peer node from the other peer nodes on the network; and the different instance of the service recognizing the peer node using the unique identifier; and the different instance of the service routing information to the peer node at the different network location.

5 99. A method comprising:

a peer node discovering an instance of a service on one of a plurality of peer nodes, wherein the one of the plurality of peer nodes is local to a network location of the peer node on a network, wherein the plurality of peer nodes each host an instance of the same service; and the peer node accessing the instance of the service;

wherein said discovering and said accessing the instance of the service are performed in accordance with one or more peer-to-peer platform protocols;

the peer node moving from the network location to a different network location; the peer node discovering the same instance of the service on the one of the plurality of peer nodes; the peer node accessing the instance of the service; and

wherein said discovering and accessing the same instance of the service are performed in accordance with the one or more peer-to-peer platform protocols;

the peer node providing a unique identifier for the peer node to the instance of the service, wherein the unique identifier distinguishes the peer node from the other peer nodes on the network; the instance of the service recognizing the peer node using the unique identifier; and the instance of the service routing information to the peer node at the different network location.

100. A method comprising:

a peer node discovering an instance of a content on one of a plurality of peer nodes, wherein the one of the plurality of peer nodes is local to a network location of the peer node on a network, wherein the plurality of peer nodes each include an instance of the same content; and the peer node accessing the instance of the content;

wherein said discovering and accessing the instance of the content are performed in accordance with one or more peer-to-peer platform protocols;

the peer node moving from the network location to a different network location; the peer node discovering a different instance of the content on a different one of the plurality of peer nodes, wherein the different one of the plurality of peer nodes is the different network location; the peer node accessing the different instance of the content;

wherein said discovering and accessing the different instance of the content are performed in accordance with the one or more peer-to-peer platform protocols.

101. An article of manufacture comprising software instructions executable to implement:

a plurality of peer nodes coupled to a network each implementing a core layer of a peer-to-peer platform, wherein the core layer comprises one or more peer-to-peer platform protocols for enabling the plurality of peer nodes to discover each other, communicate with each other, and cooperate with each other to form peer groups and share content in a peer-to-peer environment; the plurality of peer nodes each implementing a service layer comprising one or more core services each provided by one or more of the plurality of peer nodes in the peer-to-peer environment, wherein each of the one or more core services are configured to be accessed by peer nodes in the peer-to-peer environment in accordance with at least a subset of the one or more peer-to-peer platform protocols; the plurality of peer nodes each implementing an application layer comprising one or more applications each provided by one or more of the plurality of peer nodes in the peer-to-peer environment, wherein each of the one or more applications are configured to be accessed in accordance with at least one of the one or more peer-to-peer platform protocols, and wherein at least a subset of the one or more applications are each con-

figured to access at least one of the one or more core services to perform application tasks in the peer-to-peer environment in accordance with at least one of the one or more peer-to-peer platform protocols; and at least a subset of the plurality of peer nodes accessing at least a subset of the core services in accordance with at least one of the one or more peer-to-peer platform protocols to form one or more peer groups in the peer-to-peer environment.

102. The article of manufacture as recited in clause 101, wherein each of the one or more peer-to-peer platform protocols defines one or more advertisement formats for describing resources in the peer-to-peer environment, and wherein the software instructions are further executable to publish advertisements for the resources in the peer-to-peer environment, wherein the resources include one or more of the peer nodes, the peer groups, the content, the core services, other services in the service layer, the applications, pipes, and pipe endpoints, wherein the pipes are communications channels between one or more of the peer nodes, the core services, the other services and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels.

103. The article of manufacture as recited in clause 101, wherein at least a subset of the one or more peer-to-peer platform protocols defines one or more message formats configured for use in exchanging messages between the peer nodes in the peer-to-peer environment in accordance with the particular protocol.

104. The article of manufacture as recited in clause 101, wherein the one or more peer-to-peer platform protocols includes a peer discovery protocol for discovering resources in the peer-to-peer environment, wherein said discovering the resources returns one or more advertisements for the discovered resources formatted in accordance with the peer discovery protocol.

105. The article of manufacture as recited in clause 101, wherein the resources include one or more of the peer nodes, the peer groups, the content, the core services, other services in the service layer, the applications, pipes, and pipe endpoints, wherein the pipes are communications channels between one or more of the peer nodes, the core services, the other services and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels.

106. The article of manufacture as recited in clause 101, wherein the one or more peer-to-peer platform protocols includes a peer membership protocol for use by the peer nodes in applying for membership in one or more of the peer groups.

107. The article of manufacture as recited in clause 101, wherein the one or more peer-to-peer platform protocols includes a peer resolver protocol for use in sending generic search queries from one peer node to one or more other peer nodes in the peer-to-peer environment.

108. The article of manufacture as recited in clause 101, wherein the one or more peer-to-peer platform protocols include a pipe binding protocol for use in finding the physical location of a pipe endpoint and in binding to the pipe endpoint.

109. The article of manufacture as recited in clause 101, wherein the one or more peer-to-peer platform protocols include an endpoint routing protocol for enabling the peer nodes to request peer routing information to reach other peer nodes in the peer-to-peer environment, wherein pipes are communications channels between one or more of the peer nodes, the core services, other services in the service layer, and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels.

110. The article of manufacture as recited in clause 101, wherein the one or more peer-to-peer platform protocols includes a peer information protocol for enabling the peer nodes to obtain information about capabilities and status of other peer nodes in the peer-to-peer environment.

111. The article of manufacture as recited in clause 101, wherein each peer group is a collection of cooperating member peer nodes that provide a common set of services in the peer-to-peer environment.

112. The article of manufacture as recited in clause 111, wherein the one or more peer-to-peer platform protocols include a discovery protocol, wherein the common set of services includes a discovery service for use by member peer nodes in said peer group to discover advertised resources including the peer nodes and the peer groups in the peer-to-peer environment, wherein the discovery service is accessible in accordance with the discovery protocol.

113. The article of manufacture as recited in clause 111, wherein the one or more peer-to-peer platform protocols include a membership protocol, wherein the common set of services includes a membership service for use by member peer nodes in said peer group to reject or accept group membership applications, wherein the membership service is accessible in accordance with the membership protocol.

114. An article of manufacture comprising software instructions executable within a peer node to implement:

a peer node discovering an instance of a service on one of a plurality of peer nodes, wherein the one of the plurality of peer nodes is local to a network location of the peer node on a network, wherein the plurality of peer nodes each host an instance of the same service;
the peer node accessing the instance of the service;

5 wherein said discovering and said accessing the instance of the service are performed in accordance with one or more peer-to-peer platform protocols;

10 the peer node moving from the network location to a different network location;
the peer node discovering a different instance of the service on a different one of the plurality of peer nodes, wherein the different one of the plurality of peer nodes is local to the different network location;
the peer node accessing the different instance of the service; and

15 wherein said discovering and accessing the different instance of the service are performed in accordance with the one or more peer-to-peer platform protocols;

20 the peer node providing a unique identifier for the peer node to the different instance of the service, wherein the unique identifier distinguishes the peer node from the other peer nodes on the network; and
the different instance of the service recognizing the peer node using the unique identifier; and
the different instance of the service routing information to the peer node at the different network location.

115. An article of manufacture comprising software instructions executable within a peer node to implement:

25 a peer node discovering an instance of a service on one of a plurality of peer nodes, wherein the one of the plurality of peer nodes is local to a network location of the peer node on a network, wherein the plurality of peer nodes each host an instance of the same service;
the peer node accessing the instance of the service;

30 wherein said discovering and said accessing the instance of the service are performed in accordance with one or more peer-to-peer platform protocols;

35 the peer node moving from the network location to a different network location;
the peer node discovering the same instance of the service on the one of the plurality of peer nodes;
the peer node accessing the instance of the service; and

40 wherein said discovering and accessing the same instance of the service are performed in accordance with the one or more peer-to-peer platform protocols;
the peer node providing a unique identifier for the peer node to the instance of the service, wherein the unique identifier distinguishes the peer node from the other peer nodes on the network;
the instance of the service recognizing the peer node using the unique identifier; and
the instance of the service routing information to the peer node at the different network location.

116. An article of manufacture comprising software instructions executable within a peer node to implement:

45 a peer node discovering an instance of a content on one of a plurality of peer nodes, wherein the one of the plurality of peer nodes is local to a network location of the peer node on a network, wherein the plurality of peer nodes each include an instance of the same content;
the peer node accessing the instance of the content;

50 wherein said discovering and accessing the instance of the content are performed in accordance with one or more peer-to-peer platform protocols;

55 the peer node moving from the network location to a different network location;
the peer node discovering a different instance of the content on a different one of the plurality of peer nodes, wherein the different one of the plurality of peer nodes is the different network location;
the peer node accessing the different instance of the content;

wherein said discovering and accessing the different instance of the content are performed in accordance with the one or more peer-to-peer platform protocols.

[0333] Various embodiments may further include receiving, sending or storing instructions and/or data implemented in accordance with the foregoing description upon a carrier medium. Generally speaking, a carrier medium may include storage media or memory media such as magnetic or optical media, e.g., disk or CD-ROM, volatile or non-volatile media such as RAM (e.g. SDRAM, DDR SDRAM, RDRAM, SRAM, etc.), ROM, etc. as well as transmission media or signals such as electrical, electromagnetic, or digital signals, conveyed via a communication medium such as network and/or a wireless link.

[0334] Various modifications and changes may be made as would be obvious to a person skilled in the art having the benefit of this disclosure. It is intended that the invention embrace all such modifications and changes and, accordingly, the above description to be regarded in an illustrative rather than a restrictive sense.

Claims

1. A peer computing system comprising:

a plurality of peer nodes operable to couple to a network;

wherein the plurality of peer nodes are configured to implement a peer-to-peer environment on the network according to a peer-to-peer platform comprising:

a core layer comprising one or more peer-to-peer platform protocols for enabling the plurality of peer nodes to discover each other, communicate with each other, and cooperate with each other to form peer groups and share content in the peer-to-peer environment;

a service layer comprising one or more core services each provided by one or more of the plurality of peer nodes in the peer-to-peer environment, wherein at least a subset of the core services are operable to be used by the plurality of peer nodes in forming and participating in the peer groups, and wherein each of the one or more core services are configured to be accessed by the plurality of peer nodes in accordance with at least one of the one or more peer-to-peer platform protocols; and

an application layer comprising one or more applications each provided by one or more of the plurality of peer nodes in the peer-to-peer environment, wherein each of the one or more applications are configured to be accessed in accordance with at least one of the one or more peer-to-peer platform protocols, and wherein at least a subset of the one or more applications are each configured to access at least one of the one or more core services to perform application tasks in the peer-to-peer environment in accordance with at least one of the one or more peer-to-peer platform protocols.

2. The peer computing system as recited in claim 1, wherein each of the one or more peer-to-peer platform protocols defines one or more advertisement formats for describing and publishing advertisements for resources in the peer-to-peer environment, wherein the resources include one or more of the peer nodes, the peer groups, the content, the core services, other services in the service layer, the applications, pipes, and pipe endpoints, wherein the pipes are communications channels between one or more of the peer nodes, the core services, the other services and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels.

3. The peer computing system as recited in claim 1, wherein at least a subset of the one or more peer-to-peer platform protocols defines one or more message formats configured for use in exchanging messages between the peer nodes in accordance with the particular protocol.

4. The peer computing system as recited in claim 1, wherein the one or more peer-to-peer platform protocols include a discovery protocol for discovering resources in the peer-to-peer environment.

5. The peer computing system as recited in claim 4, wherein the resources include one or more of the peer nodes, the peer groups, the content, the core services, other services in the service layer, the applications, pipes, and pipe endpoints, wherein the pipes are communications channels between one or more of the peer nodes, the core services, the other services and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels.

6. The peer computing system as recited in claim 4, wherein the resources include the peer nodes, wherein the one

or more peer-to-peer platform protocols define a peer advertisement format configured for use in advertising the peer nodes in the peer-to-peer environment, wherein said discovering the peer nodes returns one or more peer advertisements for the discovered peer nodes formatted in accordance with the peer advertisement format.

- 5 7. The peer computing system as recited in claim 4 or 6, wherein the resources include the peer groups, wherein the one or more peer-to-peer platform protocols define a peer group advertisement format configured for use in advertising the peer groups in the peer-to-peer environment, wherein said discovering the peer groups returns one or more peer group advertisements formatted in accordance with the peer group advertisement format.
- 10 8. The peer computing system as recited in claim 4, 6 or 7, wherein the resources include the content, wherein the one or more peer-to-peer platform protocols define a content advertisement format configured for use in advertising the content in the peer-to-peer environment, wherein said discovering content returns one or more content advertisements formatted in accordance with the content advertisement format.
- 15 9. The peer computing system as recited in claim 4, 6, 7 or 8, wherein the resources include pipes, wherein the pipes are communications channels between one or more of the peer nodes. the core services and the applications in the peer-to-peer environment, wherein the one or more peer-to-peer platform protocols define a pipe advertisement format configured for use in advertising pipes in the peer-to-peer environment, wherein said discovering pipes returns one or more pipe advertisements formatted in accordance with the pipe advertisement format.
- 20 10. The peer computing system as recited in claim 4, 6, 7, 8 or 9, wherein the resources include pipe endpoints, wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels, wherein the one or more peer-to-peer platform protocols define an endpoint advertisement format configured for use in advertising endpoints in the peer-to-peer environment, wherein said discovering endpoints returns one or more endpoint advertisements formatted in accordance with the endpoint advertisement format.
- 25 11. The peer computing system as recited in claim 4, 6, 7, 8, 9 or 10, wherein the resources include the core services and other services provided by the peer nodes in the peer-to-peer environment, wherein the one or more peer-to-peer platform protocols define a service advertisement format configured for use in advertising the core services and the other services provided by the peer nodes in the peer-to-peer environment, wherein said discovering the core services and the other services returns one or more service advertisements formatted in accordance with the service advertisement format.
- 30 12. The peer computing system as recited in claim 1 or 4, wherein the one or more peer-to-peer platform protocols includes a peer membership protocol for use by the peer nodes in applying for membership in one or more of the peer groups.
- 35 13. The peer computing system as recited in claim 1, 4, or 12, wherein the one or more peer-to-peer platform protocols include a peer resolver protocol for use in sending generic search queries from one peer node to one or more other peer nodes in the peer-to-peer environment, wherein the search queries are sent to one or more services configured to perform searches as specified by the search queries and to generate responses to the search queries, wherein the one or more services are each hosted by one of the one or more other peer nodes.
- 40 14. The peer computing system as recited in claim 13, wherein each of the one or more services is configured to find one or more of peer, peer group, content, service, application, pipe, and pipe endpoint information in accordance with each particular search query received by the particular service handler, wherein the pipes are communications channels between one or more of the peer nodes, the core services, other services in the service layer, and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels.
- 45 15. The peer computing system as recited in claim 1, 4, 12, or 13, wherein the one or more peer-to-peer platform protocols include a pipe binding protocol for use in finding the physical location of a pipe endpoint and in binding to the pipe endpoint, wherein pipes are communications channels between one or more of the peer nodes, the core services, other services in the service layer, and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels.
- 50 55

16. The peer computing system as recited in claim 1, 4, 12, 13, or 15, wherein the one or more peer-to-peer platform protocols include an endpoint routing protocol for enabling the peer nodes to request peer routing information to reach other peer nodes.

17. The peer computing system as recited in claim 16, wherein, in said requesting peer routing information, the peer nodes are configured to use the endpoint routing protocol to send route query request messages formatted in accordance with the endpoint routing protocol to one or more router peers to request the peer routing information, wherein each of the router peers is configured to cache route information for one or more routes in the peer-to-peer environment, and wherein each of the router peers is further configured to return route information for a particular route specified by a particular route query request message if the route information for the particular route is cached by the particular router peer.

18. The peer computing system as recited in claim 17, wherein each of the router peers is further configured to forward the route query request message to other router peers if the route information for the particular route is not cached by the particular router peer.

19. The peer computing system as recited in claim 1, 4, 12, 13, 15, or 16, wherein the one or more peer-to-peer platform protocols includes a peer information protocol for enabling the peer nodes to obtain information about capabilities and status of other peer nodes in the peer-to-peer environment.

20. The peer computing system as recited in claim 1, 4, 12, 13, 15, 16, or 19, wherein each peer group is a collection of cooperating member peer nodes that provides a common set of services to the member peer nodes in the peer-to-peer environment, wherein the peer-to-peer platform protocols include a discovery protocol and a membership protocol, wherein the common set of services on at least a subset of the peer groups includes a discovery service for use by member peer nodes in said peer group to discover advertised resources including peer nodes and peer groups in the peer computing system in accordance with the discovery protocol, wherein the common set of services on at least a subset of the peer groups includes a membership service for use by member peer nodes in said peer group to reject or accept group membership applications in accordance with the membership protocol.

21. The peer computing system as recited in claim 1, wherein each of the plurality of peer nodes includes a unique identifier configured for use in distinguishing each peer node from the other peer nodes in the peer-to-peer environment.

22. A method for implementing a peer-to-peer environment on a network, the method comprising:

a plurality of peer nodes coupled to a network each implementing a core layer of a peer-to-peer platform, wherein the core layer comprises one or more peer-to-peer platform protocols for enabling the plurality of peer nodes to discover each other, communicate with each other, and cooperate with each other to form peer groups and share content in the peer-to-peer environment;

the plurality of peer nodes each implementing a service layer comprising one or more core services each provided by one or more of the plurality of peer nodes in the peer-to-peer environment, wherein each of the one or more core services are configured to be accessed by peer nodes in the peer-to-peer environment in accordance with at least a subset of the one or more peer-to-peer platform protocols;

the plurality of peer nodes each implementing an application layer comprising one or more applications each provided by one or more of the plurality of peer nodes in the peer-to-peer environment, wherein each of the one or more applications are configured to be accessed in accordance with at least one of the one or more peer-to-peer platform protocols, and wherein at least a subset of the one or more applications are each configured to access at least one of the one or more core services to perform application tasks in the peer-to-peer environment in accordance with at least one of the one or more peer-to-peer platform protocols; and

at least a subset of the plurality of peer nodes accessing at least a subset of the core services in accordance with at least one of the one or more peer-to-peer platform protocols to form one or more peer groups in the peer-to-peer environment.

23. The method as recited in claim 22, wherein the one or more peer-to-peer platform protocols include a peer membership protocol for joining or forming a peer group with other peer nodes, wherein the one or more core services include a membership service for use by the peer nodes in forming the peer groups and joining the peer groups, wherein the membership service is configured to be accessed by the peer nodes in the peer-to-peer environment in accordance with the membership protocol, the method further comprising one or more of the plurality of peer

nodes forming a peer group in the peer-to-peer environment using the membership service.

24. The method as recited in claim 23, further comprising:

5 another peer node applying for membership in the peer group using the membership service;
one or more member peer nodes of the peer group determining if the other peer node is qualified for membership in the peer group in response to the application for membership using the membership service; and
if the member peer nodes determine that the other peer node is qualified for membership in the peer group,
10 the other peer node becoming a member peer node in the peer group.

25. The method as recited in claim 22 or 23, wherein the one or more peer-to-peer platform protocols include a discovery protocol for discovering resources in the peer-to-peer environment, and wherein the one or more core services include a discovery service for use by the peer nodes to discover advertised resources in the peer-to-peer environment, wherein the discovery service is configured to be accessed by the peer nodes in the peer-to-peer environment in accordance with the discovery protocol.
15

26. The method as recited in claim 25, wherein the advertised resources include one or more of the peer nodes, the peer groups, the content, the core services, other services in the service layer, the applications, pipes, and pipe endpoints, wherein the pipes are communications channels between one or more of the peer nodes, the core services, the other services and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels.
20

27. The method as recited in claim 25, wherein the resources include the peer nodes, the method further comprising:

25 one of the plurality of peer nodes broadcasting a peer discovery message in the peer-to-peer environment using the discovery service; and
the one of the plurality of peer nodes receiving one or more response messages in response to the peer discovery message, wherein the response messages each include information about a particular peer node,
30 wherein the information is configured for use by the one of the plurality of peer nodes in establishing a connection to the particular peer node; and

wherein the peer discovery message and the one or more response messages are in a format defined by the discovery protocol, and wherein said broadcasting a peer discovery message and said receiving one or more response messages are performed in accordance with the discovery protocol.
35

28. The method as recited in claim 25 or 27, wherein the resources include the peer groups, the method further comprising:

40 one of the plurality of peer nodes broadcasting a peer group discovery message in the peer-to-peer environment using the discovery service; and
the one of the plurality of peer nodes receiving a peer group response message in response to the peer group discovery message from each of one or more of the peer groups in the peer-to-peer environment, wherein the peer group response messages each include information about a particular peer group, wherein the information is configured for use by the one of the plurality of peer nodes in joining the particular peer group; and
45

wherein the peer group discovery message and the peer group response message are in a format defined by the discovery protocol, and wherein said broadcasting a peer group discovery message and said receiving a peer group response message are performed in accordance with the discovery protocol.
50

29. The method as recited in claim 22, further comprising publishing advertisements for resources in the peer-to-peer environment using one or more advertisement formats each defined by one of the one or more peer-to-peer platform protocols, wherein the resources include one or more of the peer nodes, the peer groups, the content, the core services, other services in the service layer, the applications, pipes, and pipe endpoints, wherein the pipes are communications channels between one or more of the peer nodes, the core services, the other services and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels.
55

30. The method as recited in claim 22, further comprising two or more of the plurality of peer nodes exchanging messages in the peer-to-peer environment using one or more message formats each defined by one of the one or more peer-to-peer platform protocols.

31. The method as recited in claim 22 or 23, wherein the one or more peer-to-peer platform protocols include a discovery protocol, the method further comprising a peer node discovering resources in the peer-to-peer environment in accordance with the discovery protocol, wherein said discovering the resources comprises the peer node receiving one or more advertisements for the discovered resources formatted in accordance with the peer discovery protocol.

32. The method as recited in claim 31, wherein the resources include one or more of the peer nodes, the peer groups, the content, the core services, other services in the service layer, the applications, pipes, and pipe endpoints, wherein the pipes are communications channels between one or more of the peer nodes, the core services, the other services and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels.

33. The method as recited in claim 22, 23, 25, or 31, wherein the one or more peer-to-peer platform protocols includes a peer membership protocol, the method further comprising one of the plurality of peer nodes applying for membership in one or more of the peer groups in accordance with the peer membership protocol.

34. The method as recited in claim 22, 23, 25, 31, or 33, wherein the one or more peer-to-peer platform protocols includes a peer resolver protocol, the method further comprising one of the plurality of peer nodes sending one or more generic search queries to one or more other peer nodes in the peer-to-peer environment in accordance with the peer resolver protocol.

35. The method as recited in claim 22, 23, 25, 31, 33, or 34, wherein the one or more peer-to-peer platform protocols include a pipe binding protocol, the method further comprising:

one of the plurality of peer nodes finding the physical location of a pipe endpoint in accordance with the pipe binding protocol; and
the peer node binding to the pipe endpoint in accordance with the pipe binding protocol;

wherein pipes are communications channels between one or more of the peer nodes, the core services, other services in the service layer, and the applications in the peer-to-peer environment, and wherein the pipe endpoints are network interfaces on the peer nodes that are configured to be bound to the pipes to establish the communications channels.

36. The method as recited in claim 22, 23, 25, 31, 33, 34, or 35, wherein the one or more peer-to-peer platform protocols include an endpoint routing protocol, the method further comprising one of the plurality of peer nodes requesting peer routing information to other peer nodes in the peer-to-peer environment in accordance with the endpoint routing protocol.

37. The method as recited in claim 22, 23, 25, 31, 33, 34, 35, or 36, wherein the one or more peer-to-peer platform protocols include a peer information protocol, the method further comprising one of the plurality of peer nodes obtaining information about capabilities and status of one or more other peer nodes in the peer-to-peer environment in accordance with the peer information protocol.

38. The method as recited in claim 22, wherein each peer group is a collection of cooperating member peer nodes, further comprising each peer group providing a common set of services to the member peer nodes in the peer-to-peer environment.




FIG. 1A




FIG. 1B




FIG. 2




FIG. 3

FIG. 4

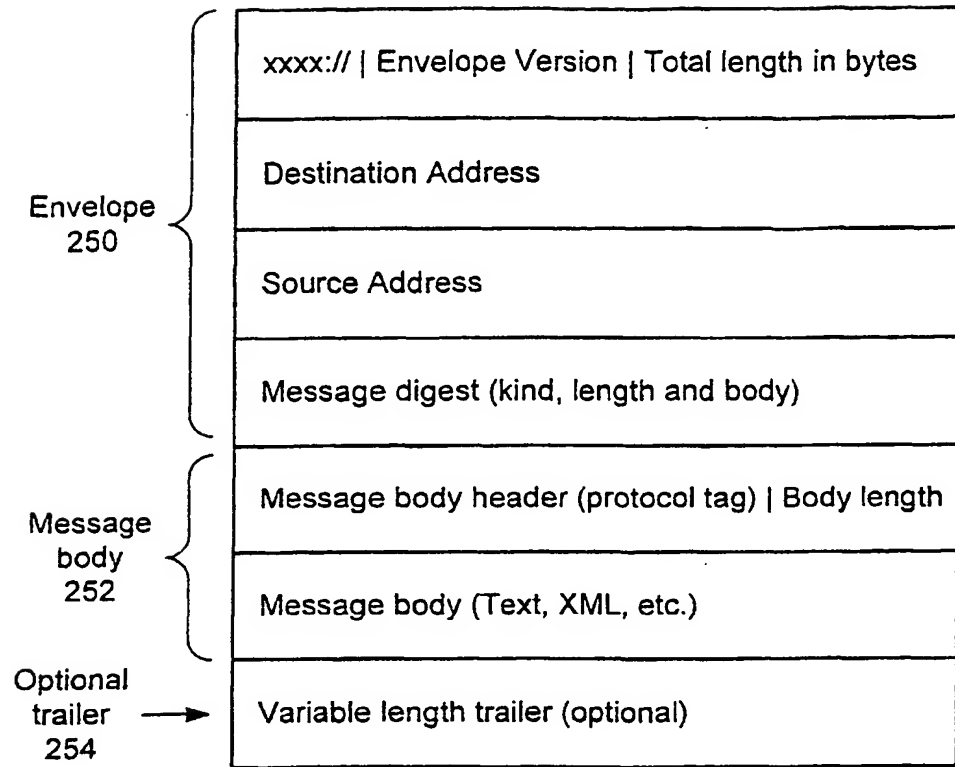


FIG. 5

Element Name (A String)	Element Value Type
Name	<String>
Keywords	<String>
Properties	<Properties>
Service	<Service Advertisement[]> (an array)
Endpoint	<Endpoint Advertisement[]> (an array)

FIG. 6

Element Name (A String)	Element Value Type
Name	<String>
Keywords	<String>
GroupID	<ID>
PeerID	<ID>
Service	<Service Advertisement[]> (an array)

FIG. 7

Element Name (A String)	Element Value Type
Name	<String>
Identifier	<ID>
Type	<String>

FIG. 8

Element Name (A String)	Element Value Type
Name	<String>
Keywords	<String>
Identifier	<ID>
Version	<String>
Pipe	<PipeAdvertisement>
Params	<String>
URI	<String>
Provider	<String>
Access Method	<Method[]> (an array)

FIG. 9

Element Name (A String)	Element Value Type
Mime-Type (Optional)	<String>
Size	Long
Encoding	<String>
ContentID	<ID>
RefID (Optional)	<ID>
Document	<Document>

FIG. 10

Element Name (A String)	Element Value Type
Name	<String>
Keywords	<String>
Address	<String>
Transport	<Transport Advertisement>

FIG. 11




FIG. 12




FIG. 13

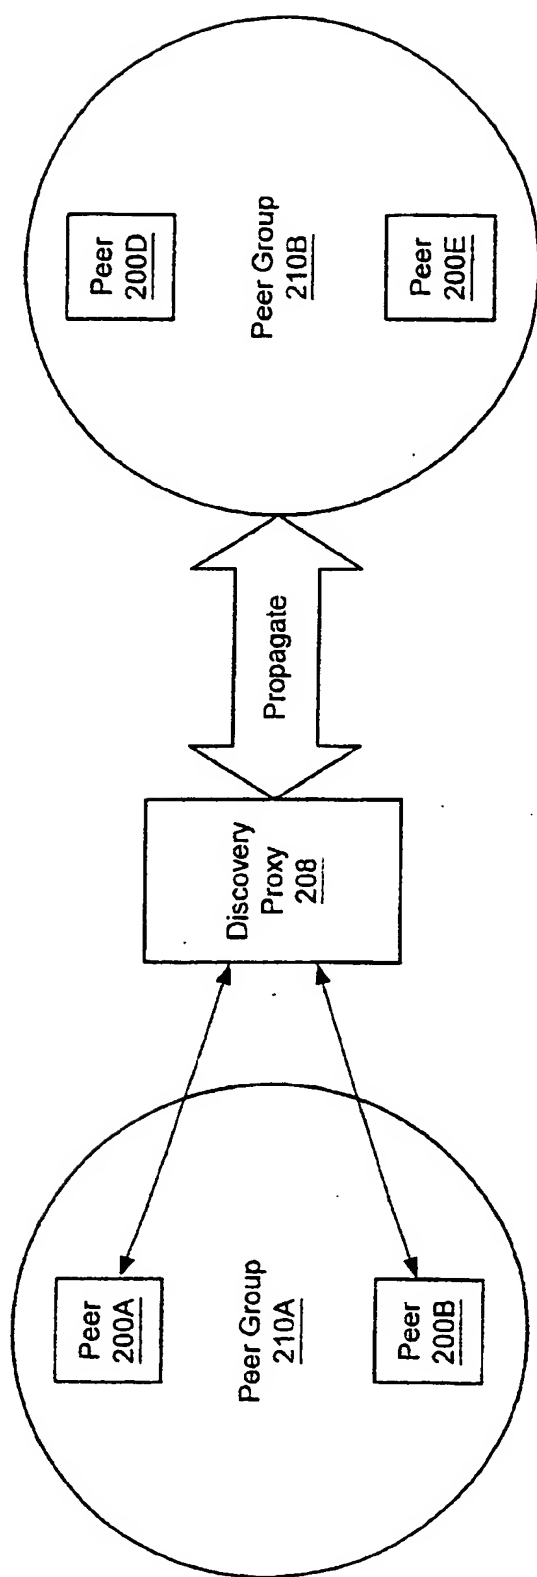


FIG. 14

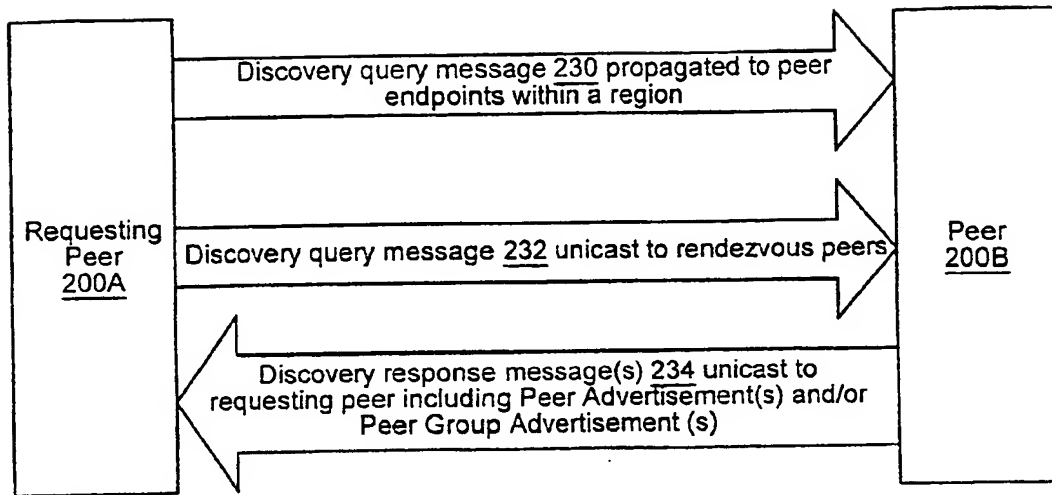


FIG. 15

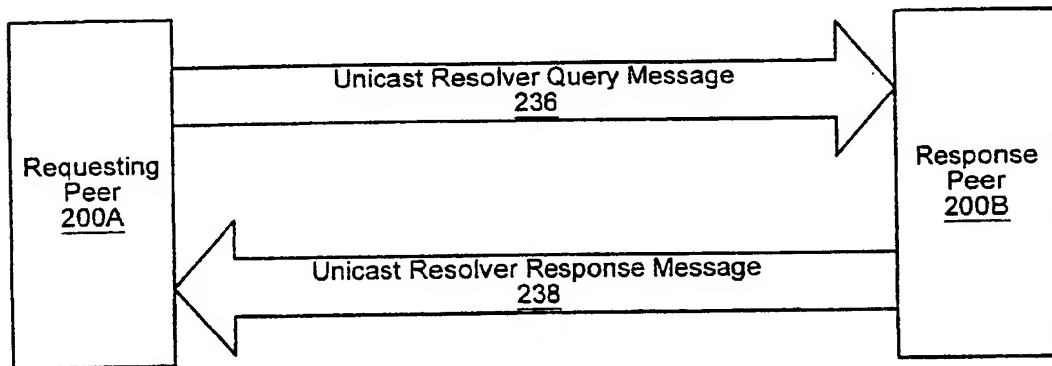


FIG. 16

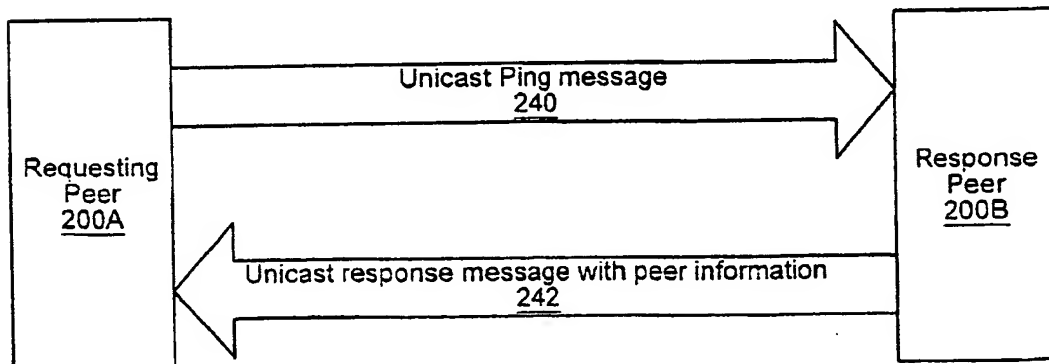


FIG. 17

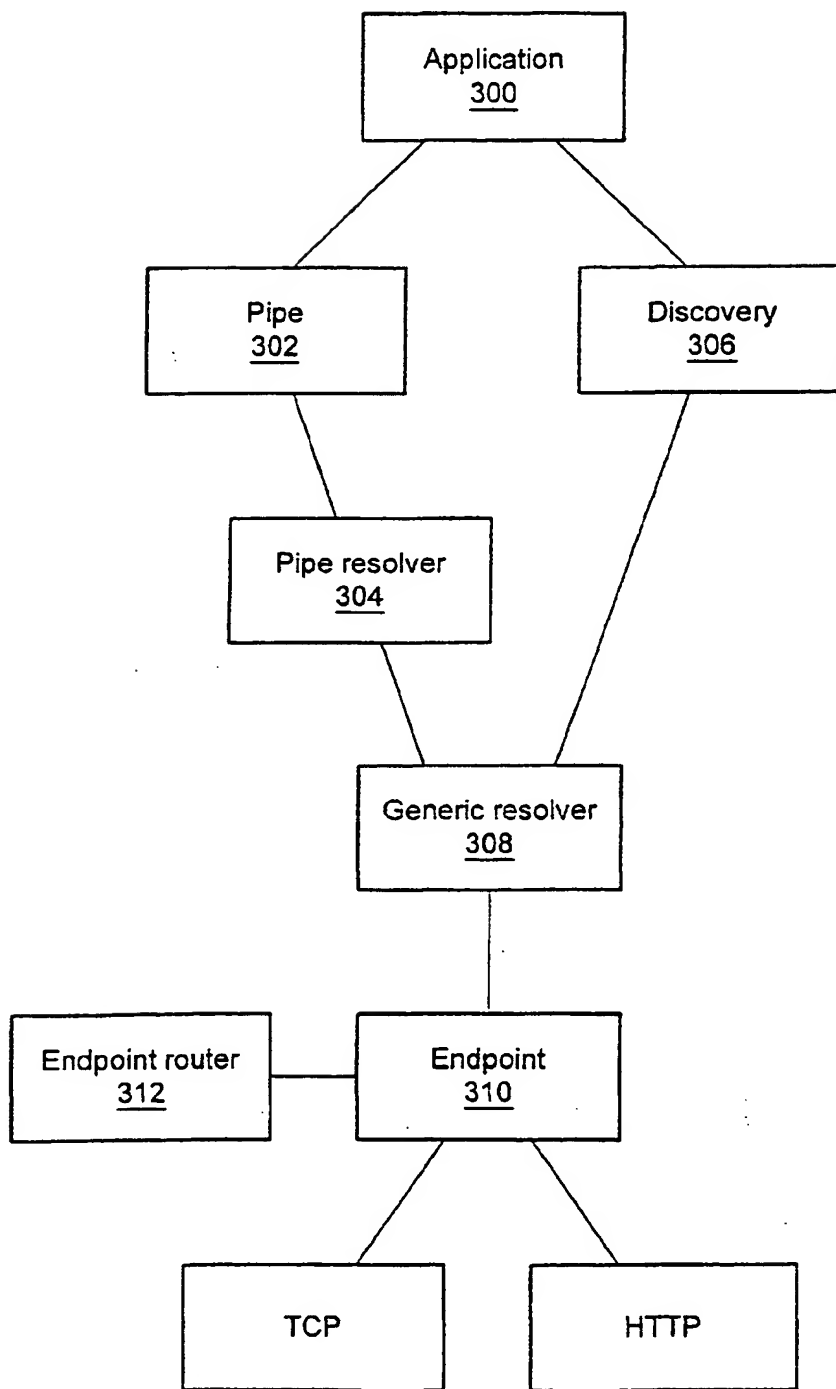


FIG. 18




FIG. 19




FIG. 20




FIG. 21

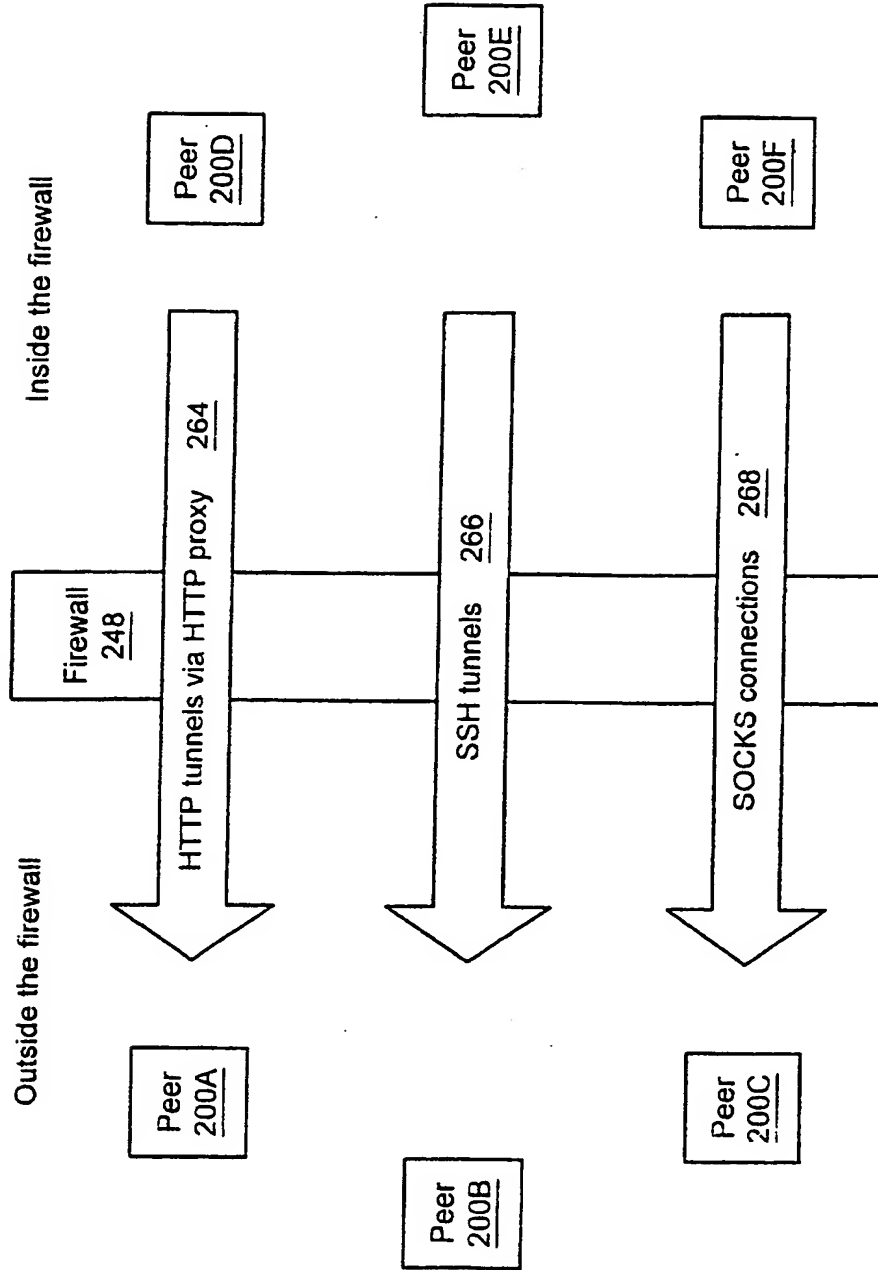


FIG. 22

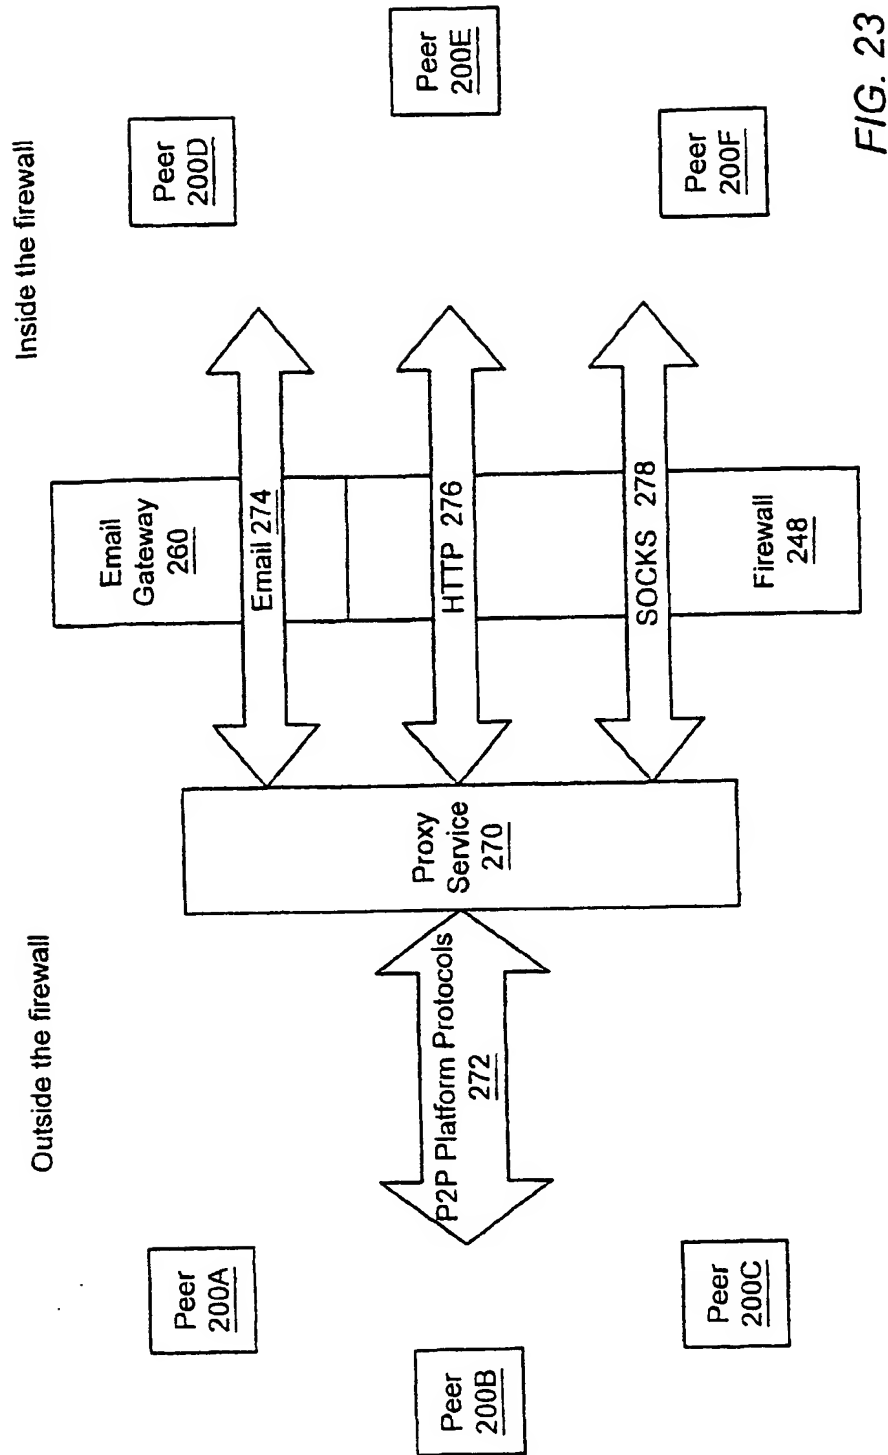


FIG. 23




FIG. 24




FIG. 25




FIG. 26




FIG. 27B




FIG. 27A




FIG. 28B




FIG. 28A

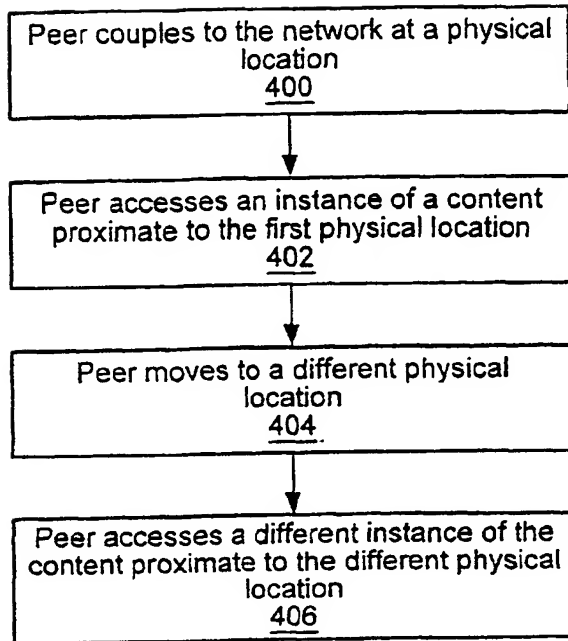


FIG. 29

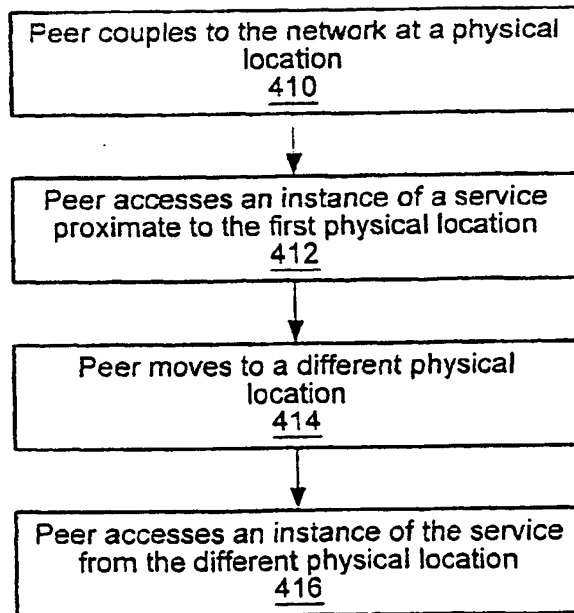


FIG. 30